arXiv:1601.06200v3 [math.OC] 15 Feb 2016

ON THE CONVERGENCE OF GMRES-ACCELERATED ADMM IN $O(\kappa^{1/4}\log\epsilon^{-1})$ ITERATIONS FOR QUADRATIC OBJECTIVES

RICHARD Y. ZHANG AND JACOB K.WHITE

Abstract. This paper investigates the generalized minimum residual method (GMRES) in its ability to accelerate the convergence of the alternating direction method-of-multipliers (ADMM). We provide evidence that ADMM-GMRES can consistently converge to an ϵ-accurate solution for a κ-conditioned problem in $O(\kappa^{1/4}\log\epsilon^{-1})$ iterations, and characterize two broad classes of problems for which the enhanced convergence is guaranteed. At the same time, we construct a class of problems that forces ADMM-GMRES to converge at the same asymptotic rate as ADMM. To demonstrate the enhanced convergence rate in practice, the accelerated method is applied to the Newton direction computation for the interior-point solution of semidefinite programs in the SDPLIB test suite.

1. Introduction

The alternating direction method-of-multipliers (ADMM) solves problems of the form

$$\begin{aligned} &\underset{x,z}{\text{minimize}} && f(x)+g(z) \\ &\text{subject to} && Ax+Bz=c \end{aligned} \tag{1.1}$$

with variables $x\in\mathbb{R}^{n_x}$ and $z\in\mathbb{R}^{n_z}$ and constant data $A\in\mathbb{R}^{n_y\times n_x}$, $B\in\mathbb{R}^{n_y\times n_z}$, and $c\in\mathbb{R}^{n_y}$. Beginning with a choice of the quadratic-penalty / step-size parameter $\beta>0$ and initial points $\{x^{(0)},z^{(0)},y^{(0)}\}$, the method generates iterates

$$\begin{aligned} \text{Local variable update: } x^{(k+1)} &= \arg\min_x f(x)+\frac{\beta}{2}\|Ax+Bz^{(k)}-c+\frac{1}{\beta}y^{(k)}\|^2, \\ \text{Global variable update: } z^{(k+1)} &= \arg\min_z g(x)+\frac{\beta}{2}\|Ax^{(k+1)}+Bz-c+\frac{1}{\beta}y^{(k)}\|^2, \\ \text{Multiplier update: } y^{(k+1)} &= y^{(k)}+\beta(Ax^{(k+1)}+Bz^{(k+1)}-c), \end{aligned}$$

that are guaranteed to converge under mild assumptions. The method finds a wide range of applications in statistics, machine learning, and related areas; cf. [6] for an extensive review.

Existing use of ADMM is mostly limited to applications where solutions of modest accuracy would be adequate. The reason is that, as a first-order method, it is subject to a fundamental trade-off between convergence speed and the smoothness of the underlying problem. A classic complexity result due to Nesterov [19, Thm. 2.1.13] asserts that, given constants $0<m\leq\ell<\infty$, no first-order method can minimize every convex objective, with gradient Lipschitz constant ℓ and strong convexity parameter m, to ϵ-accuracy with an iteration bound better than

$$O(\sqrt{\kappa}\log\epsilon^{-1}) \text{ iterations}, \tag{1.2}$$

where the condition number is defined $\kappa=\ell/m$. ADMM is known to attain (1.2) with the right choice of β and under various regularizing assumptions [14, 13, 20], but even in these cases, convergence is usually not fast enough to be competitive for high-accuracy applications.

Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139-4307. Email: ryz@mit.edu and white@mit.edu. Financial support for this work was provided in part by the Skolkovo-MIT initiative in Computational Mathematics.

This paper is motivated by a surprising observation. When ADMM is accelerated using the generalized minimum residual method (GMRES), ϵ-accurate solutions are *consistently* produced in just

$$O(\kappa^{\frac{1}{4}}\log\epsilon^{-1}) \text{ iterations.} \tag{1.3}$$

ADMM can take thousands of iterations to converge on examples where ADMM-GMRES converges in just tens of iterations. Our main results characterize two broad classes of problems, i.e. choices of objective functions and constraint matrices in (1.1), that are guaranteed to enjoy this enhanced convergence, both in theory and in practice. At the same time, we show in Section 6 that one can construct problems for which ADMM-GMRES will converge no faster than the estimate in (1.2), and so Nesterov's complexity bound is not violated. To demonstrate the effectiveness of ADMM-GMRES in applications, we use it to solve the Newton direction subproblems associated with the interior-point solution of large-scale semidefinite programs in Section 8.

1.1. **ADMM for quadratics and a surprising observation.** The general convergence properties of ADMM are most commonly analyzed using the theory of maximum monotone operators (cf. [6, Sec. 3.5] for a historical review). One may establish that ADMM converges to the solution from every initial point [12, 6], with error that scales $O(1/k)$ or $O(1/k^2)$ at the k-th iteration [18, 15], and error that scales $O(e^{-k})$ at the k-th iteration under strong convexity assumptions [9, 17, 7, 8].

The local convergence properties are best understood by modeling the objectives as quadratics, and applying classic techniques from spectral analysis. In particular, existing parameter selection and convergence rate results have mostly been derived within this context, including the variations of ADMM that attains the bound in (1.2) [4, 25, 13].

In this paper, we will restrict our attention to the quadratic-linear objectives,

$$f(x) = \frac{1}{2}xDx + p^T x, \qquad g(z) = q^T z, \tag{1.4}$$

alongside the following strong convexity assumption, which guarantees that the error will scale $O(e^{-k})$ at the k-th iteration [9].

Assumption 1 (Strong convexity)**.** The matrix D is symmetric positive definite, the matrix B has full column-rank, i.e. B^TB is invertible, and the matrix A has full row-rank, i.e. AA^T is invertible.

Defining the associated strong convexity parameter m and the gradient Lipschitz constant ℓ respectively

$$m = \lambda_{\min}(\tilde{D}), \qquad \ell = \lambda_{\max}(\tilde{D}), \qquad \tilde{D} \triangleq (AD^{-1}A^T)^{-1}, \tag{1.5}$$

the complexity lower-bound in (1.2) is attained for every problem with the parameter choice $\beta = \sqrt{m\ell}$; the leading constant in the estimate can be further improved by introducing over-relaxation [14, 13, 20].

In the context of quadratic objectives, GMRES can be applied to accelerate ADMM in a largely plug-and-play manner, to yield a consistent and significant speed-up over regular ADMM. Figure 1 makes this comparison for 1000 problems in which the A, B, D matrices are randomly generated by selecting random orthonormal bases and random singular values from a log-normal distribution.

1.2. **Main results.** For the objectives in (1.4) alongside Assumption 1, the unique solution is specified through the Karush–Kuhn–Tucker (KKT) conditions,

$$\begin{bmatrix} D & 0 & A^T \\ 0 & 0 & B^T \\ A & B & 0 \end{bmatrix} \begin{bmatrix} x^\star \\ y^\star \\ z^\star \end{bmatrix} = \begin{bmatrix} -p \\ -q \\ c \end{bmatrix} \qquad \Leftrightarrow \qquad Mu^\star = r. \tag{1.6}$$

Accordingly, ADMM reduces to linear fixed-point iterations, and GMRES convergence analysis reduces to a polynomial approximation problem over the eigenvalues of the corresponding iteration

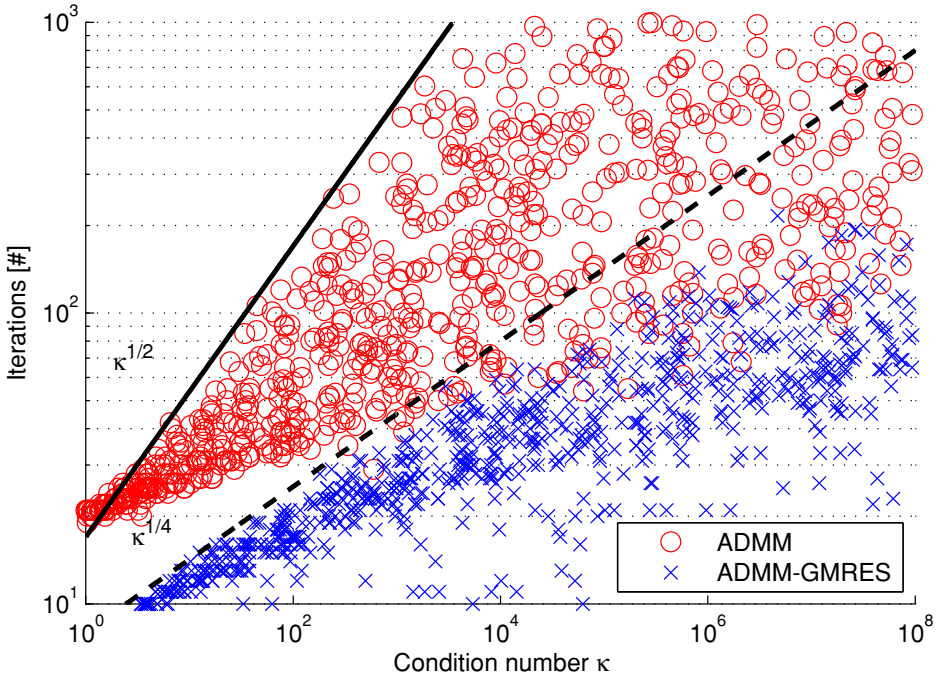

FIGURE 1. Given the same 1000 randomly-generated problems and using the same parameter choice $\beta = \sqrt{m\ell}$, ADMM (circles) converges in $O(\sqrt{\kappa})$ iterations while GMRES-accelerated ADMM (crosses) converges in $O(\kappa^{\frac{1}{4}})$ iterations. The problems have random dimensions $1 \leq n_x \leq 1000$, $1 \leq n_y \leq n_x$, $1 \leq n_z \leq n_y$. Primal-dual residual tolerance is $\epsilon = 10^{-6}$.

matrix. We will refer to the Euclidean norm of the KKT residual in dealing with notions of convergence.

Definition 1 (Residual convergence)**.** Given the initial and final iterates $u^{(0)} = [x^{(0)}; z^{(0)}; y^{(0)}]$ and $u^{(k)} = [x^{(k)}, z^{(k)}, y^{(k)}]$, we say ϵ residual convergence is achieved in k iterations if $\|Mu^{(k)} - r\| \leq \epsilon\|Mu^{(k)} - r\|$, where M and r are the KKT matrix and vector in (1.6).

We show in Section 4 that the additional square-root factor arises from a Chebyshev polynomial approximation applied to the purely-real eigenvalues of the ADMM iteration matrix. This is precisely the same mechanism that gives conjugate gradients the same square-root factor speed-up over gradient descent [16, Ch. 2-3]. However, our iteration matrix is non-normal, so our statement requires the normality qualifier in Assumption 2. The assumption is standard within this context and not particularly strong in practice; cf. Remark 16.

Theorem 2 (Dimension-based estimate)**.** *For any $A \in \mathbb{R}^{n_y \times n_x}$, $B \in \mathbb{R}^{n_y \times n_z}$, $c \in \mathbb{R}^{n_y}$, $p \in \mathbb{R}^{n_x}$, $q \in \mathbb{R}^{n_z}$, and $D \in \mathbb{R}^{n_x \times n_x}$ satisfying Assumption 1, define $\tilde{D} = (AD^{-1}A^T)^{-1}$, $m = \lambda_{\min}(\tilde{D})$, $\ell = \lambda_{\max}(\tilde{D})$, and*

$$k_0 = 2\max\{n_z, n_y - n_z\}, \qquad k_{\max} = n_y.$$

Then with the fixed choice of $\beta = \sqrt{m\ell}$, GMRES-accelerated ADMM solves (1.1) with f, g defined in (1.5) to ϵ residual convergence in

$$2 + \min\left\{k_{\max}, \quad k_0 + \left\lceil k_0\kappa^{\frac{1}{4}}\log\kappa + \kappa^{\frac{1}{4}}\log(c_1\kappa_P\kappa_X\epsilon^{-1})\right\rceil\right\} \text{ iterations,}$$

where the condition number is $\kappa = \ell/m$, the matrix normality term κ_X is defined in Assumption 2, and the scalars c_1, κ_P are defined in Lemmas 13 & 14.

Proof. The proof is provided in Section 4. □

Remark 3. The scalars c_1 and κ_P are relatively benign bookingkeeping terms, with values that grow no faster than polynomial with respect to the conditioning of the data matrices A, B, D and the choice of the parameter β. Even very large values have relatively little theoretical impact: suppose $c_1\kappa_P = 10^{12}$ at the limits of double precision; then our iteration bound for an $\epsilon = 10^{-6}$ accurate solution is only increased by a multiplicative factor of 3.

Theorem 2 guarantees the enhanced convergence rate when $k_0 \ll k_{\max}$, i.e. whenever the matrix B is very thin or almost square, subject to B^TB being nonsingular. Such problems arise, e.g. during the active-set solution of quadratic programs, which is the subject of the convergence analysis in [13].

But Theorem 2 fails to explain the computational results seen in Figure 1. In almost every one of the 1000 problems, k_0 is on the same order as $k_{\max}$, and for a few select problems, they are equal. In Section 5, we explain this puzzling phenomenon by observing that the eigenvalues with nonzero imaginary parts, which we name *complicating eigenvalues*, are often much better conditioned than the purely-real ones. This observation allows us to derive our second iteration estimate.

Theorem 4 (Coherency-based estimate)**.** *Let $A, B, D, c, p, q, m, \ell, \kappa$, and β be the same as in Theorem 2. Define δ_{lb} as in*

$$1 - \delta_{\text{lb}} = \frac{1}{2}\left[\lambda_{\max}\{\Pi_B(+H)\Pi_B\} + \lambda_{\max}\{\Pi_B^\perp(-H)\Pi_B^\perp\}\right], \tag{1.7}$$

where the two projectors are $\Pi_B = B(B^TB)^{-1}B^T$ and $\Pi_B^\perp = I - \Pi_B$, and the symmetric indefinite matrix

$$H = \frac{\sqrt{\kappa}+1}{\sqrt{\kappa}-1}\left[2A(\beta^{-1}D + A^TA)^{-1}A^T - I\right] \tag{1.8}$$

satisfies $\|H\| = 1$ by definition. Then GMRES-accelerated ADMM solves (1.1) with f, g defined in (1.5) to ϵ residual convergence in

$$2 + \left\lceil \frac{(c_2 + \delta_{\text{lb}}^{-1})(\kappa^{\frac{1}{4}}+1)}{2c_2 + (\delta_{\text{lb}}\kappa^{\frac{1}{4}})^{-1}} \log(2c_1\kappa_P\kappa_X\epsilon^{-1}) \right\rceil \text{ iterations,}$$

where $c_2 = 1/[4\log(1+\sqrt{2})]$ is an absolute constant. The matrix normality term κ_X is defined in Assumption 2, and the scalars c_1, κ_P are defined in Lemmas 13 & 14.

Proof. The proof is provided in Section 5. □

Remark 5. Loosely speaking, Theorem 4 predicts convergence to an ϵ-accurate solution in $O(\delta_{\text{lb}}^{-1}\kappa^{\frac{1}{4}}\log\epsilon^{-1})$ iterations, or $O(\sqrt{\kappa}\log\epsilon^{-1})$ iterations if $\delta_{\text{lb}}^{-1} \notin O(\kappa^{\frac{1}{4}})$, e.g. if $\delta_{\text{lb}} = 0$.

Remark 6. The exact value of δ_{lb} is driven by the *mutual coherency* between the eigenspace of $\tilde{D}$ and the column-space of B, and by the decay of eigenvalues in $\tilde{D} = (AD^{-1}A^T)^{-1}$. Consider substituting $\|H\| = 1$ into (1.7) and rewriting (1.8)

$$2\delta_{\text{lb}} \geq \|H\| - \min\left\{\|Q^THQ\|, \|P^THP\|\right\},$$
$$H = \frac{\sqrt{\kappa}+1}{\sqrt{\kappa}-1}V[(\beta^{-1}\Lambda + I)^{-1} - (\beta\Lambda^{-1} + I)^{-1}]V^T,$$

where $V\Lambda^{-1}V^T = AD^{-1}A^T$ is the eigendecomposition for $\tilde{D}^{-1}$, and Q, P are any orthogonal matrices satisfying $QQ^T = \Pi_B$ and $PP^T = \Pi_B^\perp$. If V, Q and P are mutually incoherent (e.g. in the sense defined by Donoho & Huo [10]), then we should expect the gap between $\|H\|$ and

$\min\{\|Q^THQ\|, \|P^THP\|\}$ to be nonzero in general. If the singular values of H also decay quickly, then this gap should be large, and δ_{lb} should be bounded away from zero.

For the the set of 1000 random problems shown in Figure 1, the sample mean for δ_{lb} is 0.3, and the sample minimum is 0.05. Hence, Theorem 4 sufficient to explain the enhanced convergence rate.

Finally, in Section 6, we show that there are problem constructions that force ADMM-GMRES to converge in $O(\sqrt{\kappa}\log\epsilon^{-1})$ iterations. In particular, we show that in the worst-case, the optimal polynomial that bounds the convergence of GMRES is precisely the polynomial associated with over-relaxed ADMM. In other words, ADMM-GMRES converges at the same rate as over-relaxed ADMM in the worst-case.

1.3. **Application Example: Interior-point Newton direction for SDPs.** The interior-point Newton direction problem for semidefinite programs (SDP) is a challenging environment where highly-accurate solutions are desired for dense, large-scale, and extremely ill-conditioned quadratic problems [29, 28, 27]. As progress is made and the duality gap is reduced, the system of equations becomes increasingly ill-conditioned. The condition numbers routinely grow as large as $\kappa \approx 10^8$, and the equations must be solved to a sufficiently high level of accuracy to guarantee quadratic converge in the underlying Newton's method.

The Newton direction problem provides a realistic scenario to benchmark the performance of GMRES-accelerated ADMM. In Section 8, we compare the performance of ADMM and GMRES-accelerated ADMM in their ability to recompute the Newton directions as generated by SeDuMi [26], a popular, open-source, MATLAB-based linear conic programming solver. More specifically, problems selected from the SDPLIB test suite [5] are pre-solved using SeDuMi, the Newton step problems at each interior-point step are exported, and ADMM and ADMM-GMRES are used to recompute the solution.

Our results show that ADMM-GMRES converges in $O(\kappa^{\frac{1}{4}}\log\epsilon^{-1})$ iterations over each of 1038 Newton problem considered. Many of these problems fit within the two characterizations described in this paper, and the enhanced convergence rate is explained by Theorems 2 & 4. For many others, the enhanced convergence rate is observed even when neither theorems are applicable. The result suggests further improvements to be made to the characterizations presented in this paper.

1.4. **Future work.** The theoretical and empirical results in this paper hint that ϵ-convergence to a κ-conditioned problem in $O(\kappa^{\frac{1}{4}}\log\epsilon^{-1})$ iterations may be the *average-case* behavior for ADMM-GMRES, at least in the context of the quadratic objectives of the forms in (1.4) and the associated regularity assumptions. An important next step is to make this observation more precise. One possible way to do this is to introduce a statistical framework, in order to study the probabilistic distribution of δ_{lb} in Theorem 4. As we had noted in Remark 6, the quantity is closely associated with the idea of mutual coherency, so it may be possible to adopt existing results from compressed sensing and related fields for this analysis.

2. Preliminaries

2.1. **Definitions & Notation.** Our notations are standard: upper-case Latin letters for matrices, and lower-case Latin and Greek letters for scalars and vectors. The set of real numbers is denoted $\mathbb{R}$, and the set of complex numbers is denoted $\mathbb{C}$. A complex number with zero imaginary component is said to be purely-real.

Given a matrix M, we use $\lambda_i(M)$ to refer to its i-th eigenvalue, and $\Lambda\{M\}$ to denote its set of eigenvalues, including multiplicities. If M is singular, then the notation $\Lambda_{nz}\{M\} \subseteq \Lambda\{M\}$ is used to refer to its nonzero eigenvalues. The spectral radius is the supremum of the eigenvalue moduli, and is denoted $\rho(M)$. If the eigenvalues are purely-real, then $\lambda_{\max}(M)$ refers to its most positive eigenvalue, and $\lambda_{\min}(M)$ its most negative eigenvalue. We will often refer to an eigenvalue with nonzero imaginary parts as a "complicating eigenvalue", for reasons made clear in Section 4.

Let $\|\cdot\|$ denote the l_2 vector norm, as well as the associated induced norm, also known as the spectral norm. We use $\sigma_i(M)$ to refer to the i-th largest singular value.

Finally, in describing the number of iterations to solve a κ-conditioned problem to ϵ-accuracy, we will often refer to an estimate of the form $O(\sqrt{\kappa}\log\epsilon^{-1})$ as simply $O(\sqrt{\kappa})$, with the implicit understanding that its relationship with ϵ is logarithmic.

2.2. **ADMM as linear fixed-point iterations.** Since the KKT conditions are linear, the ADMM update equations are also linear, and can be written in the form

$$u^{(k+1)} = G_{\mathrm{AD}}(\beta)u^{(k)} + b(\beta), \tag{2.1}$$

upon the vector of local, global, and multiplier variables, $u^{(k)} = [x^{(k)}; z^{(k)}; y^{(k)}]$. Fixing the value of β further reduces (2.1) to linear fixed-point iterations, whose convergence properties are entirely determined by the spectral properties of the iteration matrix, $G_{\mathrm{AD}}(\beta)$.

Following this framework, a number of previous authors have shown that, with a well-chosen value of β, ADMM is able to converge in $O(\sqrt{\kappa})$ iterations [13, 20]. In this paper, we state an explicit iteration bound, and provide a proof for completeness.

Proposition 7. *For any $A \in \mathbb{R}^{n_y\times n_x}$, $B \in \mathbb{R}^{n_y\times n_z}$, $c \in \mathbb{R}^{n_y}$, $p \in \mathbb{R}^{n_x}$, $q \in \mathbb{R}^{n_z}$, and $D \in \mathbb{R}^{n_x\times n_x}$ satisfying Assumption 1, define m, ℓ according to (1.5). Then ADMM with fixed parameter $\beta = \sqrt{m\ell}$ solves (1.1) with f, g defined in (1.4) to ϵ residual convergence in*

$$2 + \left\lceil (\sqrt{\kappa}+1)\log(c_1\kappa_M\epsilon^{-1})\right\rceil \ \textit{iterations,}$$

where the condition number is $\kappa = \ell/m$, the scalar c_1 is defined in Lemma 13, and $\kappa_M = \|M\|\|M^{-1}\|$ with M defined in (1.6).

Proof. The proof is provided in Appendix A. □

Furthermore, they show that over-relaxed ADMM can often improve convergence rates. Over-relaxation makes the substitution

$$Ax^{(k+1)} \leftarrow \omega Ax^{(k+1)} + (1-\omega)(Bz^{(k)} - c),$$

in the global variable and multiplier updates (steps 2 & 3), and setting the relaxation parameter $\omega \in (0, 2]$ to the value of ~ 1.5 [6]. If we write $G_{\mathrm{AD}}(\beta,\omega)$ as the iteration matrix associated with over-relaxed ADMM, with parameters β and ω, then the nonzero eigenvalues of $G_{\mathrm{AD}}(\beta,\omega)$ are related to those of $G_{\mathrm{AD}}(\beta)$ via the relation (cf. Corollary 31)[1]

$$\Lambda_{nz}\{G_{\mathrm{AD}}(\beta,\omega)\} = \omega\,\Lambda_{nz}\{G_{\mathrm{AD}}(\beta)\} + (1-\omega).$$

Therefore, it is reasonable to expect that further reductions in the spectral radius and the spectral norm may be achieved by a well-chosen, fixed value of ω.

Proposition 8. *Let A, B, D, c, p, q, m, ℓ, κ and β be the same as in Proposition 7. Then over-relaxed ADMM with fixed β and fixed over-relaxation parameter $\omega = 2$ solves (1.1) to ϵ residual convergence in*

$$2 + \left\lceil \frac{1}{2}(\sqrt{\kappa}+1)\log(c_1\kappa\epsilon^{-1})\right\rceil \ \textit{iterations,}$$

where the scalar c_1 is defined in Lemma 13, and $\kappa_M = \|M\|\|M^{-1}\|$ with M defined in (1.6).

Proof. The proof is provided in Appendix A. □

[1] Ghadimi *et al.* [13] also proved a version of this statement.

2.3. **Sequence acceleration via GMRES.** It is natural to ask: what further speed-ups are achievable by varying the over-relaxation parameter ω between iterations, and possibly also allowing it to become complex? For example, given the iterations, $u^{(k+1)} = Gu^{(k)} + b$, one may consider cycling through a sequence of relaxation coefficients $\omega_0, \omega_1, \omega_2, \ldots$, in a scheme sometimes known as a "high-order" over-relaxation scheme

$$u^{(k+1)} = (1-\omega_k)u^{(k)} + \omega_k G(u^{(k)} + b). \tag{2.2}$$

In fact, no choice of relaxation coefficients, fixed or varied, complex or real, can converge faster than GMRES, at least in the specific sense of step convergence.

Definition 9 (Step convergence)**.** Given the sequence $u^{(0)}$, $u^{(1)}$, …, $u^{(k)}$, $u^{(k+1)}$, we say ϵ step convergence is achieved at the k-th iteration if $\|u^{(k+1)} - u^{(k)}\| \leq \epsilon \|u^{(1)} - u^{(0)}\|$.

To sketch this characterization, we define the step-size at the k-th step as $\Delta u^{(k)} \triangleq (Gu^{(k)}+b) - u^{(k)}$ and rearrange (2.2) to reveal,

$$\Delta u^{(k)} = \left[\prod_{i=0}^{k-1}((1-\omega_i)I + \omega_i G)\right] r^{(0)} = p(G)\Delta u^{(0)}. \tag{2.3}$$

The role of each relaxation coefficient ω_i is to define a zero for the matrix polynomial $p(\cdot)$, and to rescale it to satisfy $p(1) = 1$. The optimal polynomial for step convergence is the minimizer of the step-size $\|\Delta u^{(k)}\| = \|p(G)\Delta u^{(0)}\|$. This is precisely the optimization problem solved by GMRES.

Proposition 10 (Saad & Schultz [24])**.** *Given the linear fixed-point iterations, $u = Gu + b$, and the initial point $u^{(0)}$. Let $u^{(k)}$ be the iterate generated at the k-th iteration of GMRES for the fixed-point equation $u = Gu + b$. Then the following bounds hold*

$$\frac{\|\Delta u^{(k)}\|}{\|\Delta u^{(0)}\|} \leq \min_{\substack{p\in\mathbb{P}_k \\ p(1)=1}} \frac{\|p(G)\Delta u^{(0)}\|}{\|\Delta u^{(0)}\|} \leq \min_{\substack{p\in\mathbb{P}_k \\ p(1)=1}} \|p(G)\|,$$

where $\Delta u^{(k)} = (Gu^{(k)} + b) - u^{(k)}$ and $\mathbb{P}_k$ denotes the space of order-k polynomials.

In the context of ADMM, the notion of step convergence is closely related, but not identical to that of residual convergence in Definition 1. The step-size and the residual norms are connected via a condition number, defined in Lemma 14, so convergence in one notion can be used to imply convergence in the other. Alternatively, in a right-preconditioned problem, the step-size coincides with the residual norm, so the two notions of convergence become identical. A right-preconditioned ADMM is developed in Section 7.

3. ADMM as a Block Gauss-Seidel Method

Consider the augmented Lagrangian to (1.1),

$$\mathscr{L}_\beta(x,z,y) = f(x) + g(x) + y^T(Ax + Bz - c) + \frac{\beta}{2}\|Ax + Bz - c\|^2, \tag{3.1}$$

whose saddle-point is determined by the *augmented* Karush-Kuhn-Tucker (KKT) equations

$$\begin{bmatrix} D + \beta A^TA & \beta A^TB & A^T \\ \beta B^TA & \beta B^TB & B^T \\ A & B & 0 \end{bmatrix} \begin{bmatrix} x \\ z \\ y \end{bmatrix} = \begin{bmatrix} \beta A^Tc - p \\ \beta B^Tc - q \\ c \end{bmatrix}. \tag{3.2}$$

The iterates generated by ADMM have a convenient interpretation as a block Gauss-Seidel matrix-splitting of (3.2), as in

$$\begin{bmatrix} D + \beta A^TA & 0 & 0 \\ \beta B^TA & \beta B^TB & 0 \\ A & B & -\frac{1}{\beta}I \end{bmatrix} \begin{bmatrix} x \\ z \\ y \end{bmatrix} = \begin{bmatrix} 0 & -\beta A^TB & -A^T \\ 0 & 0 & -B^T \\ 0 & 0 & -\frac{1}{\beta}I \end{bmatrix} \begin{bmatrix} x \\ z \\ y \end{bmatrix} + \begin{bmatrix} \beta A^Tc - p \\ \beta B^Tc - q \\ c \end{bmatrix}.$$

In turn, the convergence of the iterates is dictated by the *ADMM iteration matrix*, $G_{\mathrm{AD}}(\beta)$, defined

$$G_{\mathrm{AD}}(\beta) = \begin{bmatrix} D+\beta A^TA & 0 & 0 \\ \beta B^TA & \beta B^TB & 0 \\ A & B & -\frac{1}{\beta}I \end{bmatrix}^{-1} \begin{bmatrix} 0 & -\beta A^TB & -A^T \\ 0 & 0 & -B^T \\ 0 & 0 & -\frac{1}{\beta}I \end{bmatrix}. \tag{3.3}$$

The three pivot blocks correspond to the three steps of the ADMM algorithm: the local variable update is performed by inverting the block $(D+\beta A^TA)$, the global variable update is performed by inverting the block βB^TB, and the gradient ascent step scales the "constraint violation" by β and accumulates it within the variable.

Alternatively, note that the augmented KKT equations (3.2) may be obtained from the unaugmented KKT equations (1.6) by a left- multiplication with a β-dependent shear transformation matrix

$$T_{\mathrm{aug}}(\beta) = \begin{bmatrix} I & 0 & \beta A^T \\ 0 & I & \beta B^T \\ 0 & 0 & I \end{bmatrix}. \tag{3.4}$$

Therefore, ADMM may be also be viewed as a preconditioner matrix-splitting of the unaugmented KKT equations (1.6), using the *ADMM preconditioner matrix*, $P_{\mathrm{AD}}(\beta)$,

$$\begin{bmatrix} D & -\beta A^TB & A^T \\ 0 & 0 & B^T \\ A & B & -\frac{1}{\beta}I \end{bmatrix} \begin{bmatrix} x \\ z \\ y \end{bmatrix} = \begin{bmatrix} 0 & -\beta A^TB & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{\beta}I \end{bmatrix} \begin{bmatrix} x \\ z \\ y \end{bmatrix} + \begin{bmatrix} -p \\ -q \\ c \end{bmatrix}$$
$$\Leftrightarrow \quad P_{\mathrm{AD}}(\beta)u = [P_{\mathrm{AD}}(\beta) - M]u + r. \tag{3.5}$$

Note that the corresponding iteration matrix for (3.5) coincides with $G_{\mathrm{AD}}(\beta)$ defined earlier in (3.3).

3.1. **Basic spectral properties.** A key feature of ADMM is that dual feasibility is satisfied at every iteration. More specifically, when the gradient ascent step-size is chosen to coincide with the scaling of the quadratic penalty in the augmented Lagrangian, β, the following condition is satisfied at every k-th iteration with $k \geq 1$,

$$B^Ty^{(k)} + q = 0, \tag{3.6}$$

which is the second block-row of (3.5). Dual feasibility is an important clue that hints at an eigendecomposition. Returning to the interpretation of ADMM as linear fixed-point iterations, the condition (3.6) can only hold if $[0, 0, B^T]$ spans the left nullspace of the iteration matrix $G_{\mathrm{AD}}(\beta)$. Defining an orthogonal transform based on this insight reveals a block-Schur decomposition of $G_{\mathrm{AD}}(\beta)$.

Lemma 11. *Define the QR decomposition $B = QR$ with $Q \in \mathbb{R}^{n_y \times n_z}$ and $R \in \mathbb{R}^{n_z \times n_z}$, and define $P \in \mathbb{R}^{p \times (n_y - n_z)}$ as its orthogonal complement. Then defining the orthogonal matrix U and the scaling matrix $S(\beta)$,*

$$U = \left[\begin{array}{c|cc|c} I_{n_x} & 0 & 0 & 0 \\ 0 & I_{n_z} & 0 & 0 \\ 0 & 0 & P & Q \end{array}\right], \qquad S(\beta) = \left[\begin{array}{c|ccc} \beta I_{n_x} & 0 & 0 & 0 \\ \hline 0 & \beta R & 0 & 0 \\ 0 & 0 & I & 0 \\ \hline 0 & 0 & 0 & I_{n_z} \end{array}\right] \tag{3.7}$$

yields a block-Schur decomposition of $G_{\mathrm{AD}}(\beta)$

$$U^TG_{\mathrm{AD}}(\beta)U = S^{-1}(\beta) \left[\begin{array}{c|c|c} 0_{n_x} & G_{12}(\beta) & G_{13}(\beta) \\ \hline 0 & G_{22}(\beta) & G_{23}(\beta) \\ \hline 0 & 0 & 0_{n_z} \end{array}\right] S(\beta), \tag{3.8}$$

where the size $n_y \times n_y$ inner iteration matrix $G_{22}(\beta) = \frac{1}{2}I + \frac{1}{2}K(\beta)$ is defined in terms of the matrix

$$K(\beta) = \begin{bmatrix} Q^T \\ -P^T \end{bmatrix} [(\beta^{-1}\tilde{D} + I)^{-1} - (\beta\tilde{D}^{-1} + I)^{-1}] \begin{bmatrix} Q & P \end{bmatrix}, \tag{3.9}$$

and $\tilde{D} = (AD^{-1}A^T)^{-1}$.

Proof. The proof follows from routine computation and applications of the Woodbury identity; cf. Appendix B. □

We conclude that the ADMM iteration matrix, $G_{\mathrm{AD}}(\beta)$, has $n_x + n_y$ zero eigenvalues and n_y nonzero eigenvalues, all of which exactly coinciding with the eigenvalues of the inner iteration matrix $G_{22}(\beta)$.

Corollary 12 (Disk enclosure)**.** *The nonzero eigenvalues of $G_{\mathrm{AD}}(\beta)$ are enclosed within the disk on the complex plane,*

$$\mathcal{D}(\beta) = \left\{ z \in \mathbb{C} : \left| z - \frac{1}{2} \right| \leq \frac{1}{2} \max\left\{ \frac{\ell - \beta}{\ell + \beta}, \frac{\beta - m}{\beta + m} \right\} \right\}.$$

Proof. The nonzero eigenvalues of $G_{\mathrm{AD}}(\beta)$ are related to the eigenvalues of $K(\beta)$ via $\Lambda_{nz}\{G_{\mathrm{AD}}(\beta)\} = \Lambda\{G_{22}(\beta)\} = \frac{1}{2} + \frac{1}{2}\Lambda\{K(\beta)\}$. We use the spectral norm of $K(\beta)$ to enclose its eigenvalues. Obviously, $\|K(\beta)\| = \|(\beta^{-1}\tilde{D} + I)^{-1} - (\beta\tilde{D}^{-1} + I)^{-1}\|$; substituting $\ell = \lambda_{\max}(\tilde{D})$ and $m = \lambda_{\min}(\tilde{D})$ yields the desired result after some minor manipulations. □

After two ADMM iterations, the convergence behavior of ADMM becomes entirely dependent upon the inner iteration matrix.

Lemma 13. *For any β and any polynomial $p(\cdot)$, we have*

$$\|p(G_{\mathrm{AD}}(\beta))\, G_{\mathrm{AD}}^2(\beta)\| \leq c_1(\beta)\|p(G_{22}(\beta))\|,$$

where $c_1(\beta)$ is defined in terms of the matrices in Lemma 11, as in

$$c_1(\beta) = \|S(\beta)\|\|S^{-1}(\beta)\|\|G_{\mathrm{AD}}(\beta)\|^2.$$

Proof. The following is a standard identity for matrices with nilpotent blocks

$$\left[\begin{array}{c|c|c} 0_{n_x} & G_{12} & G_{13} \\ \hline 0 & G_{22} & G_{23} \\ \hline 0 & 0 & 0_{n_z} \end{array}\right]^{k+2} = \left[\begin{array}{c} G_{12} \\ \hline G_{22} \\ \hline 0 \end{array}\right] G_{22}^k \left[\begin{array}{c|c|c} 0 & G_{22} & G_{23} \end{array}\right]$$

and holds for any $k \geq 0$. Applying this identity to each monomial of $p(\cdot)$ yields the desired result. □

Setting $\beta = \sqrt{m\ell}$ minimizes the radius of the disk in Corollary 12, which in turn minimizes an estimate of the spectral norm for the inner iteration matrix, G_{22}.

3.2. **Different notions of convergence.** In the literature for ADMM, convergence is most commonly measured using the Euclidean norms of the primal and dual residuals,

$$r_{\mathrm{primal}}^{(k)} = Ax^{(k)} + Bz^{(k)} - c, \qquad r_{\mathrm{dual}}^{(k)} = \beta A^T B(x^{(k)} - x^{(k-1)}).$$

In fact, this notion of convergence is identical to that of residual convergence defined in Definition 1. To see this, consider the fixed-point equation in (3.5), and note that our definition of ϵ residual convergence is equivalent to the stopping condition $\|r_{\mathrm{primal}}^{(k)}\|^2 + \|r_{\mathrm{dual}}^{(k)}\|^2 \leq \epsilon^2$.

To convert between the notions of step convergence in Definition 9 and residual convergence in Definition 1, we use the following statement, which is self-explanatory via (3.5).

Lemma 14. *Let M, r be the KKT matrix and vector defined in (1.6), and let $u^{(k)}, u^{(k+1)}$ be two consecutive iterates generated by ADMM with parameter β. Then the step-size and the residual norms are related*

$$\kappa_P^{-1} \leq \frac{\|Mu^{(k)} - r\|}{\|u^{(k+1)} - u^{(k)}\|} \leq \kappa_P$$

where $\kappa_P = \|P_{\mathrm{AD}}(\beta)\|\|P_{\mathrm{AD}}^{-1}(\beta)\|$ and $P_{\mathrm{AD}}(\beta)$ is the ADMM preconditioner matrix as defined in (3.5).

Remark 15. Given any ϵ, define $\epsilon' = \epsilon\kappa_P^{-1}$. Lemma 14 says that ϵ' step convergence implies ϵ residual convergence, and ϵ' residual convergence implies ϵ step convergence.

4. GMRES-Accelerated ADMM Converges in $O(\kappa^{\frac{1}{4}})$ Iterations

In order to use Proposition 10 to derive convergence estimates for GMRES, it is common to reduce the polynomial norm-minimization problem into a polynomial approximation problem over the complex plane. For ADMM, this requires the following normality assumption, which is standard within this context.

Assumption 2 (κ_X is bounded)**.** The ADMM inner iteration matrix $G_{22}(\beta)$, defined in Lemma 11, is diagonalizable at $\beta = \sqrt{m\ell}$ with eigendecomposition, $G_{22}(\beta) = X\Lambda X^{-1}$. Furthermore, the condition number for the matrix-of-eigenvectors, $\kappa_X = \|X\|\|X^{-1}\|$, is bounded from above by an absolute constant.

Remark 16. Assumption 2 is implicitly evoked whenever a spectral radius estimate is used to study the convergence rate of ADMM in the Euclidean norm (cf. [16, 11] for a detail discussion). Hence it is also present in much of the existing literature on the convergence of ADMM within the quadratic setting [4, 25, 13]. The considerable predictive power of these existing results suggests that the assumption is not a particularly strong one in practice.

Accordingly, convergence analysis for ADMM-GMRES is reduced to an *eigenvalue approximation problem* over the eigenvalues of the ADMM inner iteration matrix,

$$\min_{\substack{p\in\mathbb{P}_k\\ p(1)=1}} \|p(G_{\mathrm{AD}})\| \leq c_1\kappa_X \min_{\substack{p\in\mathbb{P}_{k-2}\\ p(1)=1}} \max_{\lambda\in\Lambda\{G_{22}\}} |p(\lambda)|, \tag{4.1}$$

where the scalar c_1 is defined in Lemma 13. In this section, we establish the following upper-bound to (4.1), and use it to prove our first main result.

Lemma 17. *Let $\beta = \sqrt{m\ell}$ and $k_0 = 2\min\{n_z, n_y - n_z\}$. Then for all $k \geq k_0$, we have*

$$\min_{\substack{p\in\mathbb{P}_{k-2}\\ p(1)=1}} \max_{\lambda\in\Lambda\{G_{22}\}} |p(\lambda)| \leq \kappa^{k_0/2}\left(\frac{\kappa^{\frac{1}{4}}-1}{\kappa^{\frac{1}{4}}+1}\right)^{k-k_0}. \tag{4.2}$$

Proof. After introducing the necessarily preliminaries in Section 4.1, the proof is provided in Section 4.2. □

Proof of Theorem 2. GMRES must terminate in $k_{\max} = 2+n_y$ iterations, because the characteristic polynomial $\chi(\cdot)$ for G_{22} is at most order n_y. Applying Lemma 13 to $\chi(\cdot)$ yields

$$\|G_{\mathrm{AD}}\,\chi(G_{\mathrm{AD}})\,G_{\mathrm{AD}}\| \leq c_1\|\chi(G_{22})\| = 0, \tag{4.3}$$

and substituting (4.3) into Proposition 10 yields convergence in $k_{\max}$ iterations.

To derive the non-trivial estimate, we apply Lemma 17 to the eigenvalue approximation problem in (4.1). Converting the residual tolerance ϵ to the step-size tolerance $\epsilon' = \epsilon\kappa_P^{-1}$ via Lemma 14, taking logarithms, and applying the bound $x(1+x)^{-1} \leq \log(1+x) \leq x$ yields our desired result. □

4.1. **The Chebyshev approximation.** Upper-bounds to the eigenvalue approximation problem (4.1) can be obtained by heuristically approximating an outer enclosure over the eigenvalues. More specifically, given any outer enclosure $\mathcal{S} \subset \mathbb{C}$ satisfying $\mathcal{S} \supseteq \Lambda\{G_{22}\}$ and any heuristic order-k polynomial $q(\cdot)$ satisfying $q(1) = 1$, an upper-bound on (4.1) is established via the inequality chain

$$\min_{\substack{p\in\mathbb{P}_k\\ p(1)=1}} \max_{\lambda\in\Lambda\{G_{22}\}} |p(\lambda)| \leq \min_{\substack{p\in\mathbb{P}_k\\ p(1)=1}} \max_{z\in\mathcal{S}} |p(z)| \leq \max_{z\in\mathcal{S}} |q(z)|. \tag{4.4}$$

Clearly, any valid choice of outer enclosure $\mathcal{S}$ and heuristic polynomial q will yield upper-bounds. However, better bounds are generated by tighter enclosures and optimal approximations.

Let us begin by considering the disk-shaped enclosure $\mathcal{D} \supset \Lambda\{G_{22}\}$ in Corollary 12. The polynomial approximation problem over a disk has a well-known closed-form solution [23].

Theorem 18 (Circle approximation)**.** *Let $\mathcal{D}$ denote the disk on the complex plane, centered at $c \in \mathbb{C}$ with radius $a \in \mathbb{R}$, and let $\gamma \in \mathbb{C}\backslash\mathcal{D}$. Then the polynomial approximation problem has closed-form solution*

$$\min_{\substack{p\in\mathbb{P}_k\\ p(\gamma)=1}} \max_{z\in\mathcal{D}} |p(z)| = \left(\frac{a}{|\gamma - c|}\right)^k = \left(\frac{\kappa_D - 1}{\kappa_D + 1}\right)^k,$$

where $\kappa_D = (|\gamma - c| + a)/(|\gamma - c| - a)$ is the condition number for the disk. The minimum is attained by the monomial $p^\star(z) = (z - c)^k/|\gamma - c|^k$.

Plugging $\beta = \sqrt{m\ell}$ into Corollary 12 yields a condition number of $\kappa_D = \sqrt{\kappa}$, and applying Theorem 18 produces the upper-bound

$$\min_{\substack{p\in\mathbb{P}_k\\ p(1)=1}} \max_{\lambda\in\Lambda\{G_{22}\}} |p(\lambda)| \leq \left(1 - \frac{2}{\sqrt{\kappa} + 1}\right)^k. \tag{4.5}$$

Repeating the steps in the proof of Theorem 2 using the upper-bound (4.5) produces an iteration estimate of $O(\sqrt{\kappa})$. This is the best iteration estimate achievable using the disk enclosure $\mathcal{D}$, but the result is trivial; we have already established that ADMM (with our specified choice of β) will converge in $O(\sqrt{\kappa})$ iterations on its own, and GMRES-accelerated ADMM will always converge faster than ADMM.

A better iteration estimate requires a tighter enclosure over the eigenvalues of G_{22}. An important observation is that the eigenvalues *tend to be* purely-real, and that the number of eigenvalues with imaginary parts is correlated with the problem dimensions. To illustrate this point, two random problems with different problem dimensions are shown in Figure 2. The distribution of eigenvalues are typical; fixing n_x, n_y, n_z and randomly generating new matrices will recover essentially the same plot every time.

Purely-real eigenvalues spread over an interval is a particularly desirable structure in eigenvalue approximation problems, due to the existence of a closed-form optimal solution attributed to Chebyshev [23].

Theorem 19 (Interval Approximation)**.** *Let $\mathcal{I}$ denote the interval $[c - a, c + a]$ on the real line, and let $\gamma \in \mathbb{R}\backslash\mathcal{I}$. Then the polynomial approximation problem has closed-form solution*

$$\min_{\substack{p\in\mathbb{P}_k\\ p(\gamma)=1}} \max_{z\in\mathcal{I}} |p(z)| = \frac{1}{|T_k((\gamma - c)/a)|} \leq 2\left(\frac{\sqrt{\kappa_I} - 1}{\sqrt{\kappa_I} + 1}\right)^k,$$

where $T_k(z)$ is the degree-k Chebyshev polynomial of the first kind, and $\kappa_I = (|\gamma - c| + a)/(|\gamma - c| - a)$ is the condition number for the interval. The minimum is attained by the Chebyshev polynomial $p^\star(z) = T_k(\frac{z-c}{a})/|T_k(\frac{\gamma-c}{a})|$.

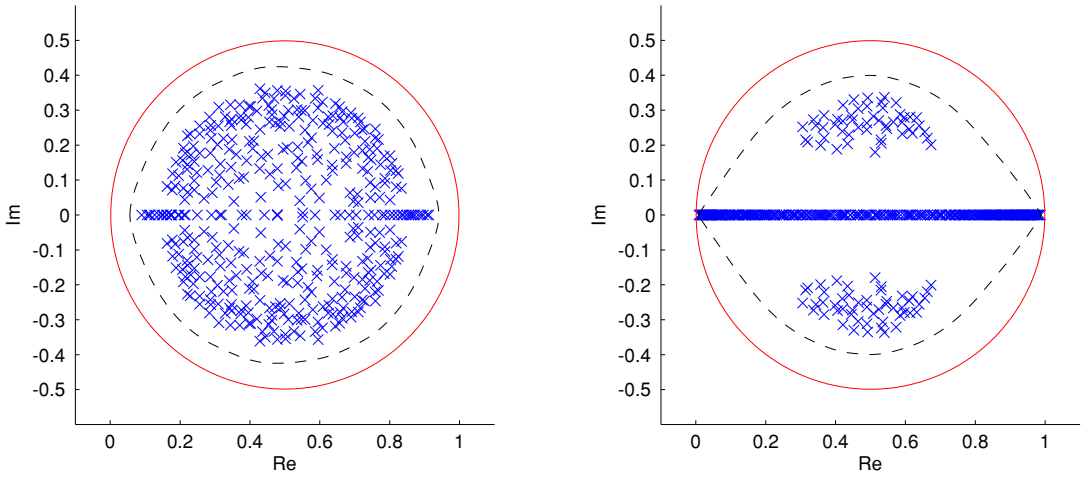

(A) (B)

FIGURE 2. Nonzero eigenvalues (crosses), field of values (dashed), and predicted disk (solid) of ADMM with $\beta = \sqrt{m\ell}$ for two randomly generated problems: (a) $n_x = 500$, $n_z = 200$, $n_y = 400$; (b) $n_x = 500$, $n_z = 50$, $n_y = 400$.

Suppose that, in addition to Corollary 12, we have *a priori* knowledge that all of our eigenvalues were purely-real. Then there exists a real interval $\mathcal{I} \supset \Lambda\{G_{22}\}$ satisfying $\kappa_I = \kappa_D = \sqrt{\kappa}$ (e.g. the projection of $\mathcal{D}$ onto the real line), and Theorem 19 assures us that there exists an order $O(\kappa^{\frac{1}{4}}\log\epsilon^{-1})$ polynomial to reduce the eigenvalue approximation error associated with $\mathcal{I}$ to below ϵ. In turn, GMRES will converge in $O(\kappa^{\frac{1}{4}})$, which is an entire square-root factor better than ADMM on its own.

4.2. **Annihilating the complicating eigenvalues.** In practice, it is very rare for all of our eigenvalues to be purely-real, i.e. there usually exists a number of eigenvalues with nonzero imaginary parts. These eigenvalues prevent the Chebyshev approximation from being directly applicable, so we refer to them as *complicating* eigenvalues. Fortunately, the number of such eigenvalues can be explicitly bounded through the problem dimensions.

Lemma 20. *The ADMM iteration matrix $G_{\mathrm{AD}}(\beta)$ has at most $2\min(n_z, n_y - n_z)$ eigenvalues with nonzero imaginary parts, counting conjugates, for every choice of $\beta > 0$.*

Proof. Recall from Lemma 11 that $G_{\mathrm{AD}}(\beta)$ has $n_x + n_z$ zero eigenvalues and n_y nonzero eigenvalues, and that the nonzero eigenvalues are shifted-and-scaled from the eigenvalues of the matrix $K(\beta)$, defined in (3.9). Define the symmetric matrix $\tilde{K}(\beta) = (\beta^{-1}\tilde{D} + I)^{-1} - (\beta\tilde{D}^{-1} + I)^{-1}$, and note that $K(\beta)$ as the following structure

$$K(\beta) = \begin{bmatrix} Q^T\tilde{K}(\beta)Q & Q^T\tilde{K}(\beta)P \\ -P^T\tilde{K}(\beta)Q & P^T\tilde{K}(\beta)P \end{bmatrix} = \begin{bmatrix} X(\beta) & Z(\beta) \\ -Z^T(\beta) & Y(\beta) \end{bmatrix}, \tag{4.6}$$

where $X = X^T \in \mathbb{R}^{n_z \times n_z}$ and $Y = Y^T \in \mathbb{R}^{(n_y - n_z)\times(n_y - n_z)}$. A matrix of this form is known as "J-symmetric", because it satisfies the symmetry condition $JK(\beta) = K^T(\beta)J$ with the matrix $J = \mathrm{blkdiag}(I_{n_z}, -I_{(n_y - n_z)})$. An immediate consequence of this J-symmetry is that $K(\beta)$ must have at most $2\min(n_z, n_y - n_z)$ eigenvalues with nonzero imaginary parts, counting conjugates [3, Prop. 2.3]. □

Suppose that the iteration matrix has k_0 complicating eigenvalues. Then, we may expend the first k_0 zeros of our polynomial (i.e. the first k_0 iterations of GMRES) to annihilating these eigenvalues.

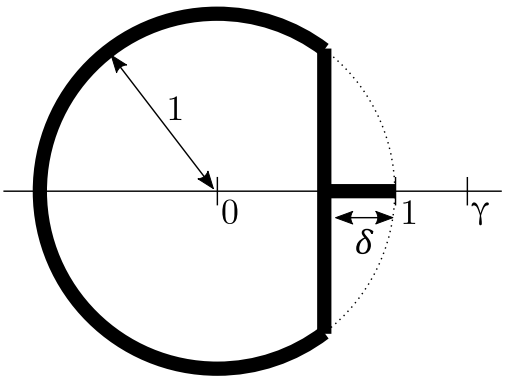

FIGURE 3. Illustration of the distribution of real and complex eigenvalues in the rescaled matrix $K/\|K\|$. The complicating eigenvalues are bounded away from the right-side of the disk by the distance δ.

Once the complicating eigenvalues are removed from consideration, the Chebyshev approximation theorem can be applied to the remaining eigenvalues, which are all purely-real. Putting these pieces together with a formal iteration bound proves the desired upper-bound on the polynomial approximation problems.

Proof of Lemma 17. Let us label the k_0 complicating eigenvalues as the set $\mathcal{C} = \{\lambda \in \Lambda\{G_{22}\} : \text{Im}\lambda \neq 0\}$, and the projection of $\mathcal{D}$ in Corollary 12 onto the real line as $\mathcal{I}$. Then the union of the two regions must satisfy $\mathcal{C} \cup \mathcal{I} \supset \Lambda\{G_{22}\}$. For a heuristic solution to the accompanying approximation problem, consider the product of the order $(k-k_0)$ Chebyshev polynomial alongside k_0 zeros placed at the complicating eigenvalues, as in

$$p(z) = \frac{T_{k-k_0}((z-c)/a)}{T_{k-k_0}((1-c)/a)} \prod_{\lambda\in\mathcal{C}} \frac{(z-\lambda)}{(1-\lambda)},$$

where $c = \frac{1}{2}$ and $a = \frac{1}{2}\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}$ are the center and the radius of $\mathcal{I}$. Clearly, $p(1) = 1$, and the polynomial is zero at every eigenvalue with a nonzero imaginary part, i.e. $\max_{\lambda\in\mathcal{C}} |p(z)| = 0$. The maximum modulus then occurs over the purely-real eigenvalues, which is bound

$$\max_{z\in\mathcal{I}} |p(z)| \leq \frac{1}{T_{k-k_0}\left(\frac{1-c}{a}\right)} \max_{z\in\mathcal{I}} \prod_{\lambda\in\mathcal{C}} \frac{|(z-\lambda)|}{|(1-\lambda)|} \leq 2(\kappa_D)^{\kappa_0} \left(\frac{\sqrt{\kappa_D}-1}{\sqrt{\kappa_D}+1}\right)^{k-k_0},$$

since $\max_{z,\lambda\in\mathcal{D}} \frac{|(z-\lambda)|}{|(1-\lambda)|} \leq \kappa_D$. Substituting $\kappa_D = \sqrt{\kappa}$ yields the desired result. □

5. A Sufficient Condition for Convergence in $O(\kappa^{\frac{1}{4}})$ Iterations

In the previous section, we showed that the additional square-root factor speed-up achieved by GMRES arises from a Chebyshev polynomial approximation of the purely-real eigenvalues of the ADMM iteration matrix. Intuitively, we would expect to see the speed-up only in cases where the vast majority of the iteration matrix eigenvalues are purely-real. Yet in our empirical results, e.g. those shown in Figure 1, we observed the speed-up over all problems, even in those that do not admit any purely-real eigenvalues.

Upon closer inspection, we find that the iteration matrix eigenvalues with non-zero imaginary parts, which we named *complicating eigenvalues*, are commonly bounded away from the constraint point $z = +1$; an illustration of this description is shown in Figure 3. Loosely speaking, the complicating eigenvalues are *better conditioned* than the purely-real eigenvalues; instead of precisely annihilating them, it may be sufficient to simply "dampen" their effect with a number of fixed-point

iterations. Then, the Chebyshev approximation can be used to approximate the remaining purely-real but poorly-conditioned eigenvalues.

Exploring this alternative eigenvalue approximation strategy, we arrive at a considerably less conservative heuristic solution to the eigenvalue approximation problem that is independent of the exact number of imaginary eigenvalues.

Lemma 21. *Let $\beta = \sqrt{m\ell}$, and define $\hat{K} \triangleq K(\beta)/\|K(\beta)\|$ as the rescaled version of matrix defined in (3.9). Define $\delta \geq 0$ as the distance from the right-most complicating eigenvalue of $\hat{K}$ to the boundary of the unit disk, as in*

$$\delta = \min\{1 - \mathrm{Re}\lambda : \lambda \in \Lambda\{\hat{K}\}, \mathrm{Im}\{\lambda\} \neq 0\}. \tag{5.1}$$

Then the polynomial approximation problem over the eigenvalues of $\hat{K}$ is bound

$$\min_{\substack{p\in\mathbb{P}_k \\ p(1)=1}} \max_{\lambda\in\Lambda\{G_{22}\}} |p(\lambda)| \leq 2\left(1 - \frac{1}{\kappa^{1/2}}\right)^{\eta}\left(1 - \frac{2}{1+\kappa^{1/4}}\right)^{\xi}, \tag{5.2}$$

where $c_2 = 1/[4\log(1+\sqrt{2})]$ and

$$\eta = \left\lceil \frac{1}{c_2\delta + 1} k \right\rceil, \qquad \xi = \left\lfloor \frac{c_2\delta}{c_2\delta + 1} k \right\rfloor. \tag{5.3}$$

Proof. The proof is provided below in Section 5.1. □

The definition of δ is illustrated in Figure 3. We will also prove the following statement, which says that the δ_{lb} as stated in Theorem 4 is a lower-bound on δ.

Proposition 22. *Let δ, β and $\hat{K}$ be as in Lemma 21, and the orthogonal matrix $[Q, P]$ according to Lemma 11. Define the scalar δ_{lb} as*

$$1 - \delta_{\mathrm{lb}} = \frac{1}{2}\left[\lambda_{\max}(Q^T\hat{K}Q) + \lambda_{\max}(-P^T\hat{K}P)\right].$$

Then $\delta_{\mathrm{lb}} \leq \delta$.

Proof. Recall from the proof of Lemma 20 that K is J-symmetric. For matrices with this structure, Benzi & Simoncini [3] used a field-of-values type argument to show that if $\lambda \in \Lambda\{K\}$ and $\mathrm{Im}\{\lambda\} \neq 0$, then

$$\frac{1}{2}\left[\lambda_{\min}(X) + \lambda_{\min}(Y)\right] \leq \mathrm{Re}\lambda \leq \frac{1}{2}\left[\lambda_{\max}(X) + \lambda_{\max}(Y)\right].$$

Substituting the definitions of δ, X, and Y results in the desired bound. □

Our second main result in this paper, i.e. the iteration bound in Theorem 4, is established by solving the eigenvalue approximation problem in (4.1) using the alternative heuristic solution in Lemma 21.

Proof of Theorem 4. Repeating the same steps as in the proof of Theorem 2, but using the bound in (5.2) in Lemma 21 yields an iteration estimate in terms of δ. Since $A(\beta^{-1}D + A^TA)^{-1}A^T = (\beta^{-1}\tilde{D} + I)^{-1}$, it is easy to verify that the matrix H in Theorem 4 satisfies $Q^THQ = Q^T\hat{K}Q$ and $P^THP = P^T\hat{K}P$. The iteration estimate is monotonously decreasing with respect to δ, and substituting $\delta_{\mathrm{lb}} \leq \delta$ from Proposition 22 yields the desired iteration estimate. □

5.1. **Damping the complicating eigenvalues.** We begin by converting the eigenvalue approximation problem over G_{22} to an equivalent problem over $\hat{K}$.

Lemma 23. *Define $\hat{K}$ as in Lemma 21, and $G_{22} \equiv G_{22}(\beta)$ as in Lemma 11 with fixed $\beta = \sqrt{m\ell}$. Then*

$$\min_{\substack{p\in\mathbb{P}_k \\ p(1)=1}} \max_{z\in\Lambda\{G_{22}\}} |p(z)| = \min_{\substack{p\in\mathbb{P}_k \\ p(\gamma)=1}} \max_{z\in\Lambda\{\hat{K}\}} |p(z)|$$

where $\gamma = (\sqrt{\kappa}+1)/(\sqrt{\kappa}-1)$.

Proof. By definition, we have $\hat{K} = K/\|K\|$ and $G_{22} = \frac{1}{2}I + \frac{1}{2}K$. Hence, the statement follows from the existence of a linear bijection between $\Lambda\{G_{22}\}\cup\{1\}$ and $\Lambda\{\hat{K}\}\cup\{\gamma\}$. □

Let δ be as defined in Lemma 21, and consider the enclosure $\mathcal{S}(\delta)\cup\mathcal{I} \supset \Lambda\{\hat{K}\}$, where

$$\mathcal{S}(\delta) = \{z\in\mathbb{C} : |z|\le 1, \mathrm{Re}\{z\}\le\delta\}, \qquad \mathcal{I} = \{z\in\mathbb{R} : |z|\le 1\}. \tag{5.4}$$

This enclosure is shown in Figure 3. The crux of our argument is a heuristic solution for the approximation problem over $\mathcal{S}(\delta)\cup\mathcal{I}$ of the form

$$p(z) = \left(\frac{z+\omega}{\gamma+\omega}\right)^{\eta} \frac{T_\xi(z)}{T_\xi(\gamma)}, \tag{5.5}$$

with polynomial orders chosen to satisfy $\eta+\xi = k$. As before, we expend ξ zeros (hence ξ iterations of GMRES) in the Chebyshev polynomial $T_\xi(z)$, in order to approximate the purely-real eigenvalues, and to produce the desired $O(\kappa^{\frac{1}{4}})$ factor in the iteration estimate. But the innovation is in the η iterations of *over-relaxation*, $(z+\omega)^\eta/(\gamma+\omega)^\eta$, which are used to resolve the complicating eigenvalues that prevent the Chebyshev approximation from being applicable. As we will soon show, the quantity δ^{-1} serves as a "condition number" over the region $\mathcal{S}(\delta)$. Reducing relative error to ϵ over every point $\lambda\in\mathcal{S}(\delta)$ requires $\eta\in O(\delta^{-1}\log\epsilon^{-1})$ iterations of over-relaxation, independent of the exact number of complicating eigenvalues and the exact value of $\gamma\ge 1$.

Lemma 24. *Define $\mathcal{S}(\delta)$ as as in (5.4), and let $\gamma\ge 1$. Then an over-relaxation solution to the complex polynomial problem yields,*

$$\min_{\substack{p\in\mathbb{P}_k \\ p(\gamma)=1}} \max_{z\in\mathcal{S}(\delta)} |p(z)| \le \left(1-\frac{\delta}{2}\right)^{k/2},$$

using the polynomial $p(z) = (1+z)^k/2^k$.

Proof. Define the over-relaxation polynomial, $p(z) = (z+\omega)^k(\gamma+\omega)^{-k}$ to satisfy $p(\gamma)=1$. For any choice of $\omega\ge 0$, the maximum is attained with the choice of $\gamma = 1$ and $z = (1-\delta)\pm j\sqrt{1-(1-\delta)^2}$, as in $\max_{z\in\mathcal{S}(\delta)}|p(z)| \le \left[1-2\delta\omega(1+\omega)^{-2}\right]^{k/2}$. Picking ω to minimize the convergence factor yields $\omega = 1$, which we use as the desired estimate. □

Remark 25. The order $\lceil 4\delta^{-1}\log(\epsilon^{-1})\rceil \in O(\delta^{-1})$ over-relaxation polynomial satisfies the conditions $\max_{z\in\mathcal{S}(\delta)}|p(z)|\le\epsilon$ and $p(\gamma)=1$.

Now, we return to the issue of the Chebyshev polynomial. Recall we had claimed in the previous section that the Chebyshev approximation is not (directly) compatible with eigenvalues with nonzero imaginary parts. The reason is that, with every increment in ξ, the maximum value of $|T_\xi(x)|$ over these eigenvalues is *increased* by a multiplicative factor.

Lemma 26. *The maximum modulus of the n-th order Chebyshev polynomial is bound within the disk centered at the origin with radius* 1*,*

$$\max_{|z|\le 1} |T_n(z)| \le T_n(\sqrt{2}) \le (1+\sqrt{2})^n, \tag{5.6}$$

and the first inequality is tight for n even.

Proof. The maximum modulus for $T_n(z)$ over the ellipse with unit focal distance and principal axis $a \geq 1$ are attained at $2n$ points along its boundary the points [23, 24]

$$z_k = a\cos\left(\frac{k\pi}{n}\right) + j\sqrt{a^2-1}\sin\left(\frac{k\pi}{n}\right) \qquad k = 1,\dots,2n.$$

The ellipse with $a = \sqrt{2}$ is the smallest to enclose the unit disk, and if n is even, then $z_{n/2}$ also lies on its boundary. The second bound follows by definition, e.g. [24, Eqn. 6.111]. □

Examining our heuristic polynomial construction in (5.5), Lemma 26 says that every increment in the Chebyshev polynomial order ξ provides a $O(\kappa^{\frac{1}{4}})$-type global error reduction, but at the cost of *locally increasing* the error about $\mathcal{S}(\delta)$. Fortunately, this error increment is a fixed constant, so may be reverted with $O(\delta^{-1})$ iterations of over-relaxation. Alternating between unit increments of ξ to reduce the global error and $O(\delta^{-1})$ increments of η to reduce the local error completes our main argument.

Proof of Lemma 21. Define the regions $\mathcal{S}(\delta)$ and $\mathcal{I}$ as in (5.4). Suppose we can establish that

$$\min_{\substack{p\in\mathbb{P}_k\\ p(\gamma)=1}} \max_{z\in\mathcal{S}(\delta)\cup\mathcal{I}} |p(z)| \leq 2\left(1-\kappa_I^{-1}\right)^{\eta}\left[1-2(\kappa_I^{1/2}+1)^{-1}\right]^{\xi}, \tag{5.7}$$

where $\kappa_I = (\gamma+1)/(\gamma-1)$ is the condition number for the interval. Then setting $\gamma = (\sqrt{\kappa}+1)/(\sqrt{\kappa}-1)$ and substituting (5.7) into Lemma 23 proves the desired statement. Beginning with $p(z)$ defined in (5.5) with $\omega = 1$, the maximum modulus over $\mathcal{S}(\delta)\cup\mathcal{I}$ is the greater of the two quantities

$$\|p(z)\|_{\mathcal{S}(\delta)} \triangleq \max_{z\in\mathcal{S}(\delta)} |p(z)| \leq \left(\frac{2\sqrt{1-\delta/2}}{\gamma+1}\right)^{\eta} \frac{(1+\sqrt{2})^{\xi}}{T_\xi(\gamma)}, \tag{5.8}$$

$$\|p(z)\|_{\mathcal{I}} \triangleq \max_{z\in\mathcal{I}} |p(z)| = \left(\frac{2}{\gamma+1}\right)^{\eta} \frac{1}{T_\xi(\gamma)} \leq 2\left(1-\frac{1}{\kappa_I}\right)^{\eta}\left(\frac{\sqrt{\kappa_I}-1}{\sqrt{\kappa_I}+1}\right)^{\xi}, \tag{5.9}$$

where the estimate in (5.8) adopts the bound in Lemma 26. We will pick the ratio η/ξ to guarantee $\|p(z)\|_{\mathcal{I}} \geq \|p(z)\|_{\mathcal{S}(\delta)}$, so that (5.9) may serve as the global error bound in (5.7). Referring (5.8) and (5.9), this means to choosing the ratio η/ξ to satisfy

$$\left(1-\frac{\delta}{2}\right)^{\eta/(2\xi)} \leq \frac{1}{1+\sqrt{2}}. \tag{5.10}$$

Viewing η/ξ as an "iteration estimate" to guarantee a relative error reduction of $1/(1+\sqrt{2})$ over $\mathcal{S}(\delta)$, we apply Lemma 24 and Remark 25 to obtain $\eta/\xi = 4\log(1+\sqrt{2})/\delta$. This choice yields (5.3) after rounding. □

5.2. **Explaining the empirical results.** Earlier in Figure 1, we presented a comparison between ADMM and GMRES-accelerated for 1000 problems. Each of these problems were constructed in the manner described in Construction 1 below. The dimension parameters n_x, n_y, n_z were uniformly sampled from $n_x \in \{1,\dots,1000\}$, $n_y \in \{1,\dots,n_x\}$, and $n_z \in \{1,\dots,n_z\}$, and the log-standard-deviation uniformly swept within the range $s \in [0,2]$.

Construction 1. Begin with nonzero positive integer parameters n_x, $n_y \leq n_x$, $n_z \leq n_y$ and positive real parameter s.

(1) Select the orthogonal matrices $U_A, U_B \in \mathbb{R}^{n_y\times n_y}$, $V_A, U_D \in \mathbb{R}^{n_x\times n_x}$, $V_B \in \mathbb{R}^{n_y\times n_z}$ i.i.d. uniformly from their respective orthogonal groups.

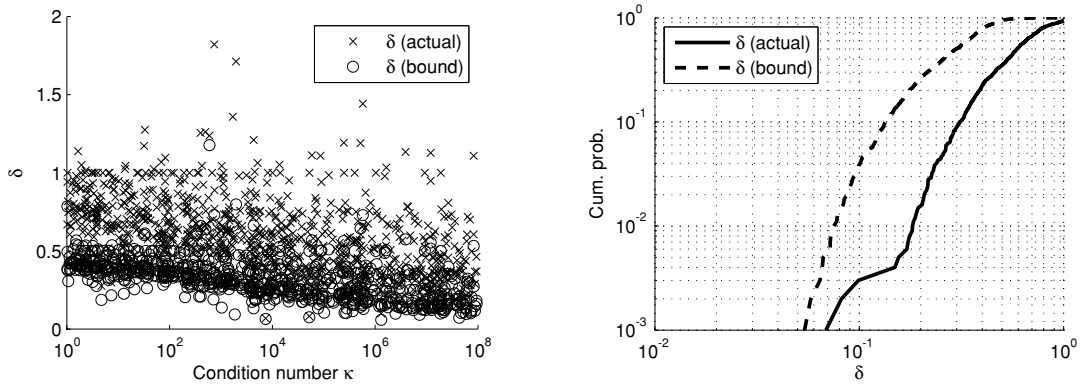

(A) (B)

FIGURE 4. Statistics for δ (cf. Theorem 4) and its lower bound (cf. Proposition 22) for 1000 randomly-generated problems: (a) scatter plot; (b) empirical CDF.

(2) Select the positive scalars $\sigma_A^{(1)},\ldots,\sigma_A^{(n_y)}$, $\sigma_B^{(1)},\ldots,\sigma_B^{(n_z)}$, and $\sigma_D^{(1)},\ldots,\sigma_D^{(n_x)}$ i.i.d. from the log-normal distribution $\sim \exp(0,s^2)$.
(3) Output the matrices $A = U_A\mathrm{diag}(\sigma_A^{(1)},\ldots,\sigma_A^{(n_y)})V_A^T$, $B = U_B\mathrm{diag}(\sigma_A^{(1)},\ldots,\sigma_A^{(n_y)})V_B^T$, and $D = U_D\mathrm{diag}(\sigma_A^{(1)},\ldots,\sigma_A^{(n_y)})U_D^T$.

To check whether Theorem 4 explains the $O(\kappa^{\frac{1}{4}})$ behavior seen in problems generated via Construction 1, we compute the value of δ for each problem considered. Figure 4a shows the distribution of δ with respect to the condition number κ for the 1000 problems previously examined in Figure 1. The smallest value of δ is 0.07, with mean and median both around 0.6. The associated cumulative probability distribution is shown in Figure 4b. An exponentially decaying probability tail can be observed. The rapid roll-off in probability tail is a signature trait for *concentration-of-measure* type results, such as the following statement.

Conjecture 27. *Fix the values of the parameters n_x, n_y, n_z, s, and select the random matrices A, B, D via Construction 1. Then there exists an absolute constant $\alpha > 0$ such that*

$$\mathbf{Pr}\left\{\delta_{\mathrm{lb}}^{-1} \geq t\,\mathbf{E}\delta^{-1}\right\} \leq e^{-\alpha(t-1)} \qquad \forall t > 1.$$

Assuming that $\mathbf{E}\delta^{-1}$ does not become too large, the conjecture suggests that GMRES converges in $O(\kappa^{\frac{1}{4}})$ iterations almost surely, because it would be extremely unlikely for a random problem generated by Construction 1 to produce a value of δ^{-1} so large as to invalidate Theorem 4.

6. Worst-case Convergence in $O(\sqrt{\kappa})$ Iterations

When the number of complicating eigenvalues k_0 approaches the total number of eigenvalues $k_{\max}$ in Theorem 2, and when the ratio δ approaches zero too quickly in Theorem 4, both of our characterizations break down, and fail to predict convergence in $O(\kappa^{\frac{1}{4}})$ iterations. Using only the disk characterization in Corollary 12, the optimal polynomial given by Theorem 18 is

$$p^\star(z) = \left(\frac{z-\frac{1}{2}}{1-\frac{1}{2}}\right)^k = (2z-1)^k, \tag{6.1}$$

which is precisely over-relaxation with the relaxation parameter set to $\omega = 2$. Its associated iteration estimate coincides with over-relaxed ADMM in Proposition 8 up to a logarithmic factor, and supersedes the estimate in Theorem 4 with $\delta = 0$.

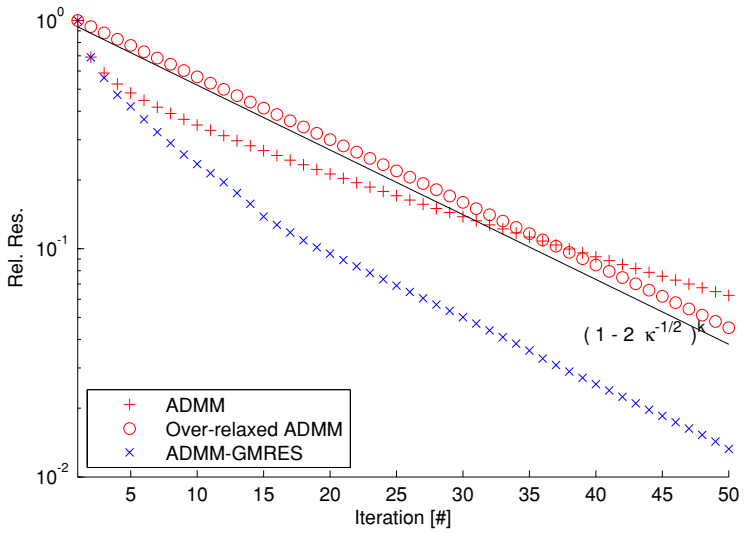

FIGURE 5. The problem construction in Section 6 places eigenvalues in a circle, so ADMM-GMRES convergences at the same asymptotic rate as over-relaxed ADMM with $\omega = 2$.

In fact, this result is sharp. In this section, we provide a class of problems (satisfying $k_0 = k_{\max}$ and $\delta = 0$) that forces GMRES to converge at the same rate as over-relaxed ADMM, attaining convergence in $O(\sqrt{\kappa})$ iterations. Since GMRES must converge in at most $O(\sqrt{\kappa})$ iterations to match the convergence of the usual ADMM, these problems represent the worst-case scenario for GMRES.

6.1. **An explicit worst-case construction.** Consider the following choice of A, B, D for any choice of condition number κ.

Construction 2. Begin with nonzero positive integer parameter n, and nonzero real positive parameter $\kappa \geq 1$.

(1) Output A as the size-$2n$ identity matrix, and $D = \text{blkdiag}(\kappa^{-\frac{1}{2}} I_n, \kappa^{\frac{1}{2}} I_n)$.
(2) Output $B \in \mathbb{R}^{2n\times n}$ in terms of the size-n diagonal matrix Θ,

$$B = \begin{bmatrix} \cos\Theta \\ \sin\Theta \end{bmatrix}, \qquad \Theta = \frac{\pi}{4n}\text{diag}\left(1, 3, 5, \ldots, 2n-1\right).$$

We will show that the optimal polynomial for this particular construction is the over-relaxation polynomial in (6.1).

Proposition 28. *Chose A, B, D according to Construction 2, and set $\beta = \sqrt{m\ell}$. Then for all $k < 2n$,*

$$\min_{\substack{p\in\mathbb{P}_k \\ p(1)=1}} \|p(G_{22})\| = \left(\frac{\kappa^{\frac{1}{2}} - 1}{\kappa^{\frac{1}{2}} + 1}\right)^k,$$

and the optimal polynomial is $p^\star(z) = (2z-1)^k$.

Despite the optimality stated in Proposition 28, we should still expect to GMRES converge in fewer iterations than over-relaxed ADMM, as shown in Figure 5. The discrepancy arises because GMRES optimizes convergence for a *specific* initial vector, while the polynomial in (6.1) optimizes convergence for the *worst-case* initial vector. More specifically, GMRES minimizes $\|p(G_{22})\Delta u^{(0)}\|/\|\Delta u^{(0)}\|$ (cf. Proposition 10), while over-relaxed ADMM minimizes $\|p(G_{22})\|$. The

latter upper-bound is indeed sharp, and there does exist an initial vector for which the iterates generated by GMRES and over-relaxed ADMM will coincide. However, in a "typical" instance, GMRES will tend to converge in slightly fewer iterations [16, 11, Ch. 3].

The key step to establish Proposition 28 is to show that the eigenvalues of $\hat{K}$ lie equally spaced on a circle. Under these circumstances, the optimal polynomial is known explicitly to be the over-relaxation polynomial.

Claim 29. The matrix $\hat{K} \triangleq K(\beta)/\|K(\beta\|$ is normal, and its eigenvalues are the (rotated) $2n$-th roots of unity,

$$\Lambda = w^{\frac{1}{2}}\,\mathrm{diag}(1, w, w^2, \ldots, w^{2n-1}), \qquad w = e^{j\pi/n}. \tag{6.2}$$

Proof. The construction yields $\tilde{D} = (AD^{-1}A^T)^{-1} = D$ and $\beta = \sqrt{m\ell} = 1$. Let us define $J = \mathrm{blkdiag}(I_n, -I_n)$ and $W = \begin{bmatrix} Q & P \end{bmatrix}$. We begin by noting that

$$(\tilde{D}+I)^{-1} - (\tilde{D}^{-1}+I)^{-1} = \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\begin{bmatrix} I_n & 0 \\ 0 & -I_n \end{bmatrix} = \|K\|J.$$

Hence, the renormalized matrix is simply $\hat{K} = JW^TJW$. Since Q is defined to span the column space of B, and P to be its orthogonal complement, there must exist orthogonal matrices $\hat{Q}$, $\hat{P}$ satisfying

$$W = \begin{bmatrix} \cos\Theta & -\sin\Theta \\ \sin\Theta & \cos\Theta \end{bmatrix}\begin{bmatrix} \hat{Q} & 0 \\ 0 & \hat{P} \end{bmatrix}.$$

Using the fact that J commutes with $\mathrm{blkdiag}(\hat{Q}, \hat{P})$, we have

$$\begin{aligned} \hat{K} &= \begin{bmatrix} \hat{Q}^T & 0 \\ 0 & \hat{P}^T \end{bmatrix} J \begin{bmatrix} \cos\Theta & \sin\Theta \\ -\sin\Theta & \cos\Theta \end{bmatrix} J \begin{bmatrix} \cos\Theta & -\sin\Theta \\ \sin\Theta & \cos\Theta \end{bmatrix}\begin{bmatrix} \hat{Q} & 0 \\ 0 & \hat{P} \end{bmatrix} \\ &= \begin{bmatrix} \hat{Q}^T & 0 \\ 0 & \hat{P}^T \end{bmatrix}\begin{bmatrix} \cos 2\Theta & -\sin 2\Theta \\ \sin 2\Theta & \cos 2\Theta \end{bmatrix}\begin{bmatrix} \hat{Q} & 0 \\ 0 & \hat{P} \end{bmatrix}. \end{aligned} \tag{6.3}$$

As shown, $\hat{K}$ is unitarily similar to $\mathrm{blkdiag}\{\exp(2j\Theta), \exp(-2j\Theta)\}$. Hence it is normal, and its eigenvalues are given as Λ in (6.2). □

Proof of Proposition 28. Since $\hat{K}$ is normal, $G_{22} = \frac{1}{2}I + \frac{1}{2}\|K\|\hat{K}$ is also normal, and the following holds exactly,

$$\|p(G_{22})\| = \max_{\lambda\in\Lambda\{G_{22}\}} |p(\lambda)|.$$

As we have shown the eigenvalues of $\hat{K}$ to be the roots of unity, the eigenvalues $\Lambda\{G_{22}\}$ are $2n$ points equally spaced along the disk $\mathcal{D}$ from Corollary 12. The closed-form solution is known via the discrete analog of Theorem 18 (cf. [23]) to be $p^\star(z) = (z-\frac{1}{2})^k/|1-\frac{1}{2}|^k = (2z-1)^k$. □

6.2. **General trends that causes slow-down to $O(\sqrt{\kappa})$.** The above is a rather artificial construction that spreads the eigenvalues of $\hat{K}$ evenly spaced around the unit circle, which causes GMRES to lose its extra square-root factor and converge in $O(\sqrt{\kappa})$ iterations. But this "slow-down" can be induced under more general circumstances. For example, consider making a change-of-basis in D, and replacing the carefully chosen B matrix with any random matrix of the same size as follows.

Construction 3. Begin with nonzero positive integer parameter n, and nonzero real positive parameter $\kappa \geq 1$.

(1) Output A as the size-$2n$ identity matrix.

(2) Select $U_D \in \mathbb{R}^{2n\times 2n}$ uniformly from the size-$2n$ orthogonal group, and output

$$D = U_D \begin{bmatrix} \kappa^{-\frac{1}{2}} I_n & 0 \\ 0 & \kappa^{+\frac{1}{2}} I_n \end{bmatrix} U_D^T.$$

(3) Select $B \in \mathbb{R}^{2n\times n}$ randomly, independent of U_D, from any desired distribution.

With overwhelming probability, GMRES will also require $O(\sqrt{\kappa})$ iterations to solve a problem from Construction 3. The cause is almost identical to that of Construction 2: the eigenvalues of $\hat{K}$ are "smeared" around the unit circle, and the optimal polynomial for the unit circle is over-relaxation. In this latter construction, the eigenvalue distribution around the circle is not completely uniform, but this is not necessary for the $O(\kappa^{\frac{1}{4}})$ factor to be lost.

As a general trend, we expect GMRES to converge in $O(\sqrt{\kappa})$ iterations whenever the complicating eigenvalues of $\hat{K}$ spread into a ring of radius ≈ 1, since the optimal polynomial under these conditions is over-relaxation. But if this distribution is not sufficiently "spread-out", then GMRES will be able to exploit clustering, and the $O(\kappa^{\frac{1}{4}})$ factor can often be recovered. A precise characterization of "boundary problems" at the edge of the slow-down remains unknown, and is the subject of future work.

7. Left- and Right- Preconditioning

In the presence of finite precision, GMRES may prematurely terminate before finding a sufficiently accurate solution, in the sense of the KKT residual norm in Definition 1. The reason for this is the factor of κ_P that links step convergence with KKT residual convergence, described earlier in Lemma 14. GMRES must terminate once the step-size reaches machine precision, but this does not guarantee that the KKT residual norm has also reach machine precision, particularly once the value of κ_P grows to be large.

Recall from (3.5) in Section 3, that the fixed-point equation for ADMM may be written as

$$u = P^{-1}(\beta)[P(\beta) - M]u + b \qquad \Leftrightarrow \qquad P^{-1}(\beta)[Mu - r] = 0, \tag{7.1}$$

where $P(\beta)$ is the ADMM preconditioner matrix. This is simply a left-preconditioned version of the KKT equation $Mu = r$. Note that its own residual vector coincides with our usual definition of the step vector. Its residual norm then coincides with our usual definition of the step-size.

Instead, consider the right-preconditioned system,

$$\hat{u} = [P(\beta) - M]P^{-1}(\beta)\hat{u} + r \qquad \Leftrightarrow \qquad MP^{-1}(\beta)\hat{u} - r = 0, \tag{7.2}$$

whose solution satisfies $P(\beta)u = \hat{u}$. The eigenvalues are clearly the same for both (7.1) and (7.2), and their norms are related by the factor κ_P. However, note that residual vector for the right-preconditioned system *coincides* with the residual of the original KKT system. In other words, consider the iterations

$$\hat{u}^{(k+1)} = [P(\beta) - M]P^{-1}(\beta)\hat{u}^{(k)} + r, \tag{7.3}$$

and observe that its step vector, $\Delta\hat{u}^{(k)} = \hat{u}^{(k+1)} - \hat{u}^{(k)}$, is identical to the residual vector in Definition 1. So if GMRES applied to (7.2) terminates at machine precision, then we can be sure that the solution vector, $u = P^{-1}(\beta)\hat{u}$, also satisfies the KKT equation up to machine precision.

8. Application Example: Interior-Point Newton Direction for Semidefinite Programs

In this section, we compare the performance of ADMM and GMRES-accelerated ADMM in their ability to recompute the Newton steps as generated by SeDuMi [26] over 80 problems in the

SDPLIB suite [5]. All 80 problems are semidefinite programs (SDP), which are given in linear conic programming form as the primal-dual pair

$$\begin{aligned} &\text{Primal: minimize } c^T y \text{ subject to } B^T y = d, \quad y \in \mathcal{K}, \\ &\text{Dual: maximize } d^T z \text{ subject to } Bz + x = c, \quad x \in \mathcal{K}^*, \end{aligned} \tag{8.1}$$

with $\mathcal{K} = \mathcal{K}^*$ set to a Cartesian product of semidefinite cones. The Newton direction at each iteration is defined as the solution to

$$\begin{bmatrix} D & 0 & I \\ 0 & 0 & B^T \\ I & B & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta z \\ \Delta y \end{bmatrix} = \begin{bmatrix} r_c \\ r_p \\ r_d \end{bmatrix}, \tag{8.2}$$

where B is the (usually) *sparse* data matrix, D is the (usually) *dense* scaling matrix specific to the interior-point method, and r_c, r_p, r_d are residual vectors. SeDuMi uses primal-dual Nesterov–Todd (NT) scaling, alongside the Mehrotra predictor-corrector method to minimize the number of interior-point steps taken. For our purposes, this means that two Newton direction subproblems are solved per interior-point step: a predictor step and a corrector step. The data matrix B is fixed for all steps, the scaling matrix D varies between steps but remains fixed between the two subproblems, and the residual vectors r_c, r_p and r_d varies between steps and also between subproblems.

The actual semidefinite programs considered are the 80 problems in the SDPLIB test suite with less than 700 constraints, i.e. with $n_z \leq 700$. For each of these problems, the predictor and corrector step subproblems are extracted, alongside the solutions computed by SeDuMi. As expected, the scaling matrix D became progressively ill-conditioned as the interior-point method progressed, with the condition number κ reaching machine precision within approximately 10 steps. In total, 1038 Newton direction problems had $\kappa \leq 10^8$, all of which we would expect GMRES to solve on the order of hundreds of iterations.

Note that the application of ADMM to semidefinite programs is not new. A number of previous authors have applied ADMM directly (i.e. without the second-order interior-point layer) to solve semidefinite programs [30, 21]. The use of ADMM to solve the Newton step problem was considered in [1, 22]. Finally, the Newton system (8.2) can be viewed as a generic saddle-point problem, and in this greater context, preconditioned Krylov-subspace iterations are widespread; cf. [2] for an extensive review.

8.1. **Efficient implementation of ADMM.** A direct solution of the Newton direction via dense Cholesky factorization of the Schur complement has cubic complexity of $O(n_y^{3/2} n_z + n_y n_z^2 + n_z^3)$. In order for ADMM to be competitive with the direct method, the parameter choice β and each iteration update, written

$$\begin{bmatrix} \Delta x^{(i+1)} \\ \Delta z^{(i+1)} \\ \Delta y^{(i+1)} \end{bmatrix} = \begin{bmatrix} D + \beta I & 0 & 0 \\ \beta B^T & \beta B^T B & 0 \\ I & B & -\frac{1}{\beta} I \end{bmatrix}^{-1} \left(\begin{bmatrix} 0 & -\beta B & -I \\ 0 & 0 & -B^T \\ 0 & 0 & -\frac{1}{\beta} I \end{bmatrix} \begin{bmatrix} \Delta x^{(i)} \\ \Delta z^{(i)} \\ \Delta y^{(i)} \end{bmatrix} + \begin{bmatrix} r_c \\ r_p \\ r_d \end{bmatrix} \right),$$

must all be available at less than cubic complexity. To be more specific, we require the following three key ingredients at subcubic complexity:

(1) Access to the extremal eigenvalues of $D \equiv \tilde{D}$, written $m = \lambda_{\min}(D)$, $\ell = \lambda_{\max}(D)$, in order to determine the parameter choice $\beta = \sqrt{m\ell}$.
(2) Matrix-vector products with $(D + \beta I)^{-1}$ for the local variable update.
(3) Matrix-vector products with $(B^T B)^{-1}$ for the global variable update.

In fact, within the context of the Newton direction for interior-point methods, the first two ingredients are always always available at subquadratic complexity of $O(n_y^{3/2})$. To be more specific, assuming that the primal-scaling, dual-scaling, or primal-dual Nesterov-Todd (NT) scaling is used

for the underlying interior-point method, the matrix D will be provided in the Kronecker product form

$$D = \text{blkdiag}(W_1 \otimes W_1, W_2 \otimes W_2, \ldots, W_N \otimes W_N).$$

Each of these constituent matrices, $W_1, \ldots, W_N$, is size-$O(n_y^{1/2})$ symmetric positive definite, so diagonalizing them requires just $O(n_y^{3/2})$ operations. Let us write each $W_i = V_i \Lambda_i V_i^T$, and define the matrices

$$\begin{aligned}\Lambda &= \text{blkdiag}(\Lambda_1 \otimes \Lambda_1, \Lambda_2 \otimes \Lambda_2, \ldots, \Lambda_N \otimes \Lambda_N),\\ V_r &= \text{blkdiag}(V_1 \otimes I, V_2 \otimes I, \ldots, V_N \otimes I),\\ V_l &= \text{blkdiag}(I \otimes V_1, I \otimes V_2, \ldots, I \otimes V_N).\end{aligned}$$

Then the eigendecomposition for D is given,

$$D = V_r V_l \Lambda V_l^T V_r^T, \qquad (D+\beta I)^\alpha = V_r V_l (\Lambda + \beta I)^\alpha V_l^T V_r^T.$$

Matrix-vector products with V_r and V_l (as well as their transposes) each requires $O(n_y^{3/2})$ operations. Furthermore, the extremal eigenvalues of D are easily accessible once each Λ_i is known

$$m = \min_{i\in\{1,\ldots,N\}} \|\Lambda_i^{-1}\|^{-2}, \qquad \ell = \min_{i\in\{1,\ldots,N\}} \|\Lambda_i\|^2.$$

Access to the third ingredient, i.e. an efficient matrix-vector product with $(B^TB)^{-1}$, is much more problem specific. In semidefinite relaxations of combinatorial optimization problems, the matrix-vector product may be available at linear complexity of $O(n_z)$, because B^TB is often either diagonal or the identity matrix by construction. But for more general problems in which the matrix B is fully dense, computing its Cholesky factorization requires cubic complexity of $O(n_y n_z^2)$. We note for these problems that the B matrix does not vary between interior-point steps, so once the Cholesky factorization of B^TB is precomputed once, the factorization may be reused for all subsequent steps at quadratic complexity of $O(n_z^2)$.

8.2. **Newton steps for problems in the SDPLIB suite.** The 1038 Newton direction problems with $\kappa \le 10^8$ are solved using standard ADMM and ADMM-GMRES with right-preconditioning, both to an accuracy of 10^{-6} relative residual for the system (8.2). The maximum number of iterations for all three methods is capped at 500. The number of iterations to convergence are shown in Figure 6a.

The results validate the $O(\sqrt{\kappa})$ figure expected of ADMM, and the $O(\kappa^{\frac{1}{4}})$ figure expected of GMRES. In fact, the multiplicative constants associated with each appear to be very similar to the results shown earlier in Figure 1. However, examining the values of δ_{lb} in Figure 6b, we find that Theorem 4 is unable to explain the $O(\kappa^{\frac{1}{4}})$ convergence rate in all of the problems. Two such cases are given as examples in the discussion below. This result suggests future work to improve the characterization in Theorem 4, in order to explain the enhanced convergence rate in these other problems.

8.3. **Example 1: 5th predictor step of control3.** In this first example, we consider the fifth predictor step for the SDPLIP problem "control3", with 1126 decision variables and 136 constraints, and a condition number of $\kappa = 4.505 \times 10^7$. With $n_y = 1126$ and $n_z = 136$, the value of $k_0 = 136$ does not allow Theorem 2 to predict an enhanced convergence rate. Computing δ yields a bound of 6.11×10^{-5} and an actual value of 1.19×10^{-3}, so Theorem 4 predicts convergence in around $O(\kappa^{\frac{1}{2}})$ iterations. However, the actual convergence performance is closer to $O(\kappa^{\frac{1}{4}})$. Explicitly forming and computing all 1126 eigenvalues of G_{22} reveals that all of the eigenvalues are close to being purely real. Hence, we may consider Theorem 2 to take into effect with $k_0 = 0$.

8.4. **Example 2: 7th predictor step of hinf14.** In this second example, we consider the seventh predictor step for the SDPLIP problem "hinf14", with 421 decision variables and 73 constraints, and a condition number of $\kappa = 3.114 \times 10^8$. With $n_y = 421$ and $n_z = 73$, the value of $k_0 = 73$ does not allow Theorem 2 to predict an enhanced convergence rate. Computing δ yields a bound of 1.636×10^{-4} and an actual value of 4.301×10^{-4}, so Theorem 4 also indicates convergence in $O(\kappa^{\frac{1}{2}})$ iterations. Again, the actual convergence performance is closer to $O(\kappa^{\frac{1}{4}})$. Explicitly forming and computing all 421 eigenvalues of G_{22} shows that the conservative value of δ is a "false" estimate, because the associated complicating eigenvalue has an imaginary part that is only on the order of 10^{-5}. Eliminating that eigenvalue, e.g. with a single zero of the polynomial, the value of δ is increased to ≈ 0.1. In this latter case, we may consider Theorem 4 to take into effect after a certain number of iterations to improve the value of δ.

Appendix A. Convergence rates for ADMM and over-relaxed ADMM

For completeness, we present short proofs for Propositions 7 & 8, i.e. that both ADMM and over-relaxed ADMM converge in $O(\sqrt{\kappa}\log\epsilon^{-1})$ iterations. We begin by proving the case for ADMM, using the much of the existing machinery throughout this paper. We then make some minor adjustments to this machinery in order to extend the proof to over-relaxed ADMM.

Proof of Proposition 7. Recall that ADMM generates iterates via the ADMM iteration matrix $G_{\mathrm{AD}}(\beta)$ defined in (3.3),

$$u^{(k+1)} = G_{\mathrm{AD}}(\beta)u^{(k)} + b, \tag{A.1}$$

where $u^{(k)} = [x^{(k)}; z^{(k)}; y^{(k)}]$, and that ϵ residual convergence is achieved if

$$\|Mu^{(k)} - r\| \leq \epsilon\|Mu^{(0)} - r\|,$$

where M and r are the KKT matrix and residual in (1.6).

Denote the fixed-point of the sequence as $u^\star$. By definition, $u^\star$ satisfies both $u^\star = G_{\mathrm{AD}}(\beta)u^\star + b$ as well as $Mu^\star = r$. Subtracting $u^\star$ from each side of (A.1) yields the error update equation and the residual update equation

$$(u^{(k+1)} - u^\star) = G_{\mathrm{AD}}(\beta)(u^{(k)} - u^\star) \tag{A.2}$$

$$Mu^{(k+1)} - r = [MG_{\mathrm{AD}}(\beta)M^{-1}](Mu^{(k)} - r). \tag{A.3}$$

Inductively repeating this argument the relative residual at the k-th iteration

$$\frac{\|Mu^{(k)} - r\|}{\|Mu^{(0)} - r\|} \leq \|MG^k_{\mathrm{AD}}(\beta)M^{-1}\| \leq \kappa_M\|G^k_{\mathrm{AD}}(\beta)\|, \tag{A.4}$$

where $\kappa_M = \|M\|\|M^{-1}\|$. To resolve this last bound we use $\|G^k_{\mathrm{AD}}(\beta)\| \leq c_1\|G^{k-2}_{22}(\beta)\|$ from Lemma 13, and substitute $\beta = \sqrt{m\ell}$ into Lemma 11 to produce $\|G_{22}(\beta)\|$ to yield

$$(\text{A.4}) \leq c_1\kappa_M\|G_{22}(\beta)\|^{k-2} \leq c_1\kappa_M\left(1 - \frac{1}{1+\kappa^{1/2}}\right)^{k-2}. \tag{A.5}$$

Solving for ϵ using $x(1+x)^{-1} \leq \log(1+x) \leq x$ yields the desired iteration estimate. □

To proceed, we will write $G_{\mathrm{AD}}(\beta,\omega)$ as the iteration matrix associated with over-relaxed ADMM with parameter ω. To extend the same proof to over-relaxed ADMM, we will need the following Lemma, which very slightly extends Lemma 11.

Lemma 30. *Define $Q, P, R, U, S(\beta)$ as in Lemma 11. Then the matrix U produces a block-Schur decomposition of $G_{\mathrm{AD}}(\beta,\omega)$*

$$U^T G_{\mathrm{AD}}(\beta,\omega)U = S^{-1}(\beta)\left[\begin{array}{c|c|c} 0_{n_x} & G_{12}(\beta,\omega) & G_{13}(\beta,\omega) \\ \hline 0 & G_{22}(\beta,\omega) & G_{23}(\beta,\omega) \\ \hline 0 & 0 & 0_{n_z} \end{array}\right] S(\beta), \tag{A.6}$$

where the size $n_y \times n_y$ inner iteration matrix is defined in terms of $\tilde{D} = (AD^{-1}A^T)^{-1}$

$$G_{22}(\beta) = \left(1-\frac{\omega}{2}\right) I + \frac{\omega}{2}\begin{bmatrix} Q^T \\ -P^T \end{bmatrix}\left[(\beta^{-1}\tilde{D}+I)^{-1} - (\beta\tilde{D}^{-1}+I)^{-1}\right]\begin{bmatrix} Q & P \end{bmatrix}. \tag{A.7}$$

Corollary 31. *Let $\Lambda_{nz}\{\cdot\}$ denote the nonzero eigenvalues of a matrix. Then*

$$\Lambda_{nz}\{G_{\mathrm{AD}}(\beta,\omega)\} = \omega\,\Lambda_{nz}\{G_{\mathrm{AD}}(\beta)\} + (1-\omega).$$

Proof of Proposition 8. Repeating the same argument as the proof of Proposition 7, but replacing all instances of Lemma 11 with Lemma 30, we arrive at

$$\frac{\|Mu^{(k)} - r\|}{\|Mu^{(0)} - r\|} \le c_1\kappa_M \|G_{22}(\beta,\omega)\|^{k-2} \le c_1\kappa_M\left(1 - \frac{\omega}{1+\kappa^{1/2}}\right)^{k-2}, \tag{A.8}$$

where $\kappa_M = \|M\|\|M^{-1}\|$ and $c_1 \le \|S^{-1}(\beta)G_{\mathrm{AD}}(\beta,\omega)\|\|G_{\mathrm{AD}}(\beta,\omega)S(\beta)\|$. Solving for ϵ using $x(1+x)^{-1} \le \log(1+x) \le x$ yields the desired iteration estimate. □

Appendix B. Proof of Lemmas 11 & 30

We will only consider the over-relaxed case Lemma 30, since the case of regular ADMM in Lemma 11 is an immediate corollary obtained by fixing $\omega = 1$. Substituting the definitions in Section 2 yields the iteration matrix

$$G_{\mathrm{AD}}(\beta,\omega) = \begin{bmatrix} D+\beta A^TA & 0 & 0 \\ \omega\beta B^TA & \beta B^TB & 0 \\ \omega A & B & -\frac{1}{\beta}I \end{bmatrix}^{-1}\begin{bmatrix} 0 & -\beta A^TB & -A^T \\ 0 & (\omega-1)\beta B^TB & -B^T \\ 0 & (\omega-1)B & -\frac{1}{\beta}I \end{bmatrix}. \tag{B.1}$$

Multiplying by U, U^T, then factoring from the right by $S(\beta)$ yields

$$U^T G_{\mathrm{AD}}(\beta)U = S^{-1}(\beta)\left[\begin{array}{c|c|c} 0 & G_{12}(\beta,\omega) & G_{13}(\beta,\omega) \\ \hline 0 & G_{22}(\beta,\omega) & G_{23}(\beta,\omega) \\ \hline 0 & 0 & 0 \end{array}\right] S(\beta), \tag{B.2}$$

where the sub-blocks are respectively

$$G_{12}(\beta,\omega) = -(\beta^{-1}D + A^TA)^{-1}A^T\begin{bmatrix} Q & P \end{bmatrix}, \tag{B.3}$$

$$G_{13}(\beta,\omega) = -(\beta^{-1}D + A^TA)^{-1}A^TQ, \tag{B.4}$$

$$G_{22}(\beta,\omega) = \begin{bmatrix} (1-\omega)I & 0 \\ 0 & I \end{bmatrix} + \omega\begin{bmatrix} Q^T \\ -P^T \end{bmatrix} A(\beta^{-1}D + A^TA)^{-1}A^T\begin{bmatrix} Q & P \end{bmatrix}, \tag{B.5}$$

$$G_{23}(\beta,\omega) = \begin{bmatrix} -B^TQ \\ 0 \end{bmatrix} + \omega\begin{bmatrix} Q^T \\ -P^T \end{bmatrix} A(\beta^{-1}D + A^TA)^{-1}A^TQ. \tag{B.6}$$

We will make the following claim to simplify the expressions

Claim 32. Let AA^T and D be invertible. Then

$$A(\beta^{-1}D + A^TA)^{-1}A^T = (\beta^{-1}\tilde{D} + I)^{-1}$$

where $\tilde{D} = (AD^{-1}A^T)^{-1}$.

Proof. The factorization arises from two applications of the Woodbury formula,

$$\begin{aligned} A(\beta^{-1}D + A^TA)^{-1}A^T &= \beta AD^{-1}A^T + \beta AD^{-1}A^T(\beta AD^{-1}A^T + I)^{-1}\beta AD^{-1}A^T \\ &= \beta\tilde{D}^{-1} + \beta\tilde{D}^{-1}(\beta\tilde{D}^{-1} + I)^{-1}\beta\tilde{D}^{-1} \\ &= (\beta^{-1}\tilde{D} + I)^{-1}. \end{aligned}$$

Finally, the Woodbury formula gives $I - (\beta^{-1}\tilde{D} + I)^{-1} = (\beta\tilde{D}^{-1} + I)^{-1}$. □

Applying the above then yields

$$G_{22}(\beta,\omega) = \left(1 - \frac{\omega}{2}\right) I + \frac{\omega}{2}\begin{bmatrix} Q^T \\ -P^T \end{bmatrix}\left[(\beta^{-1}\tilde{D} + I)^{-1} - (\beta\tilde{D}^{-1} + I)^{-1}\right]\begin{bmatrix} Q & P \end{bmatrix}, \tag{B.7}$$

which completes our proof of Lemmas 11 & 30.

References

[1] M. Annergren, S. K. Pakazad, A. Hansson, and B. Wahlberg, *A distributed primal-dual interior-point method for loosely coupled problems using ADMM*, arXiv preprint arXiv:1406.2192, (2014).

[2] M. Benzi, G. H. Golub, and J. Liesen, *Numerical solution of saddle point problems*, Acta numerica, 14 (2005), pp. 1–137.

[3] M. Benzi and V. Simoncini, *On the eigenvalues of a class of saddle point matrices*, Numerische Mathematik, 103 (2006), pp. 173–196.

[4] D. Boley, *Local linear convergence of the alternating direction method of multipliers on quadratic or linear programs*, SIAM Journal on Optimization, 23 (2013), pp. 2183–2207.

[5] B. Borchers, *SDPLIB 1.2, a library of semidefinite programming test problems*, Optimization Methods and Software, 11 (1999), pp. 683–690.

[6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Foundations and Trends® in Machine Learning, 3 (2011), pp. 1–122.

[7] D. Davis and W. Yin, *Convergence rate analysis of several splitting schemes*, arXiv preprint arXiv:1406.4834, (2014).

[8] ———, *Faster convergence rates of relaxed Peaceman-Rachford and ADMM under regularity assumptions*, arXiv preprint arXiv:1407.5210, (2014).

[9] W. Deng and W. Yin, *On the global and linear convergence of the generalized alternating direction method of multipliers*, Journal of Scientific Computing, (2012), pp. 1–28.

[10] D. L. Donoho and X. Huo, *Uncertainty principles and ideal atomic decomposition*, Information Theory, IEEE Transactions on, 47 (2001), pp. 2845–2862.

[11] T. A. Driscoll, K.-C. Toh, and L. N. Trefethen, *From potential theory to matrix iterations in six steps*, SIAM review, 40 (1998), pp. 547–578.

[12] J. Eckstein and D. P. Bertsekas, *On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators*, Mathematical Programming, 55 (1992), pp. 293–318.

[13] E. Ghadimi, A. Teixeira, I. Shames, and M. Johansson, *Optimal parameter selection for the alternating direction method of multipliers (ADMM): quadratic problems*, Automatic Control, IEEE Transactions on, 60 (2015), pp. 644–658.

[14] P. Giselsson and S. Boyd, *Diagonal scaling in Douglas-Rachford splitting and ADMM*, in Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on, IEEE, 2014, pp. 5033–5039.

[15] T. Goldstein, B. O'Donoghue, S. Setzer, and R. Baraniuk, *Fast alternating direction optimization methods*, SIAM Journal on Imaging Sciences, 7 (2014), pp. 1588–1623.

[16] A. Greenbaum, *Iterative methods for solving linear systems*, vol. 17, Siam, 1997.

[17] D. Han and X. Yuan, *Local linear convergence of the alternating direction method of multipliers for quadratic programs*, SIAM Journal on numerical analysis, 51 (2013), pp. 3446–3457.

[18] B. He and X. Yuan, *On the o(1/n) convergence rate of the Douglas-Rachford alternating direction method*, SIAM Journal on Numerical Analysis, 50 (2012), pp. 700–709.

[19] Y. Nesterov, *Introductory lectures on convex optimization*, vol. 87, Springer Science & Business Media, 2004.

[20] R. Nishihara, L. Lessard, B. Recht, A. Packard, and M. I. Jordan, *A general analysis of the convergence of ADMM*, arXiv preprint arXiv:1502.02009, (2015).

[21] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd, *Operator splitting for conic optimization via homogeneous self-dual embedding*, arXiv preprint arXiv:1312.3039, (2013).

[22] S. K. Pakazad, A. Hansson, and M. S. Andersen, *Distributed primal-dual interior-point methods for solving loosely coupled problems using message passing*, arXiv preprint arXiv:1502.06384, (2015).
[23] T. J. Rivlin, *The Chebyshev Polynomials: From Approximation Theory to Algebra and Number Theory*, John Wiley & Sons, 1974.
[24] Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on scientific and statistical computing, 7 (1986), pp. 856–869.
[25] W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin, *On the linear convergence of the ADMM in decentralized consensus optimization*, Signal Processing, IEEE Transactions on, 62 (2014), pp. 1750–1761.
[26] J. F. Sturm, *Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones*, Optimization methods and software, 11 (1999), pp. 625–653.
[27] K.-C. Toh, *Solving large scale semidefinite programs via an iterative solver on the augmented systems*, SIAM Journal on Optimization, 14 (2004), pp. 670–698.
[28] K.-C. Toh and M. Kojima, *Solving some large scale semidefinite programs via the conjugate residual method*, SIAM Journal on Optimization, 12 (2002), pp. 669–691.
[29] L. Vandenberghe and S. Boyd, *A primal-dual potential reduction method for problems involving matrix inequalities*, Mathematical Programming, 69 (1995), pp. 205–236.
[30] Z. Wen, D. Goldfarb, and W. Yin, *Alternating direction augmented lagrangian methods for semidefinite programming*, Mathematical Programming Computation, 2 (2010), pp. 203–230.

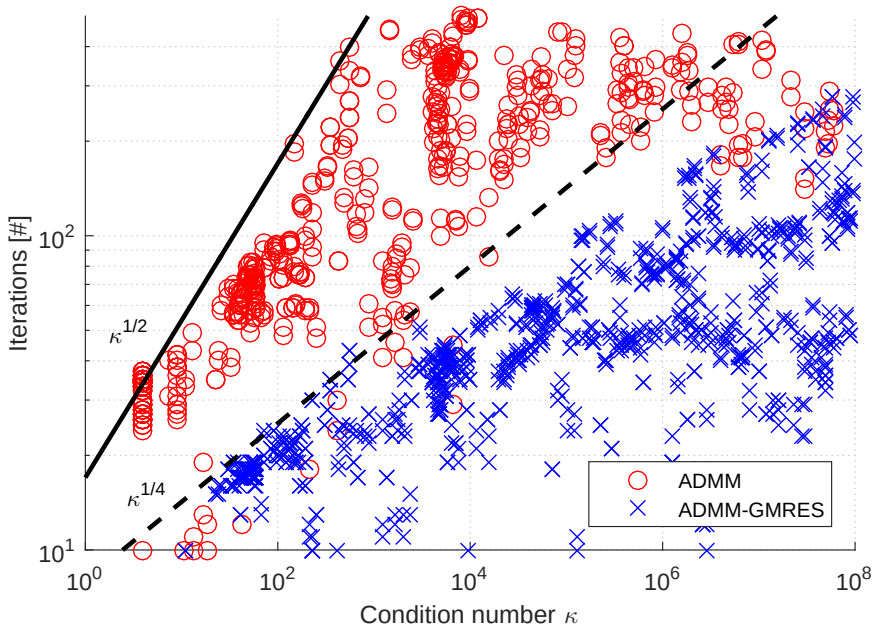

(A)

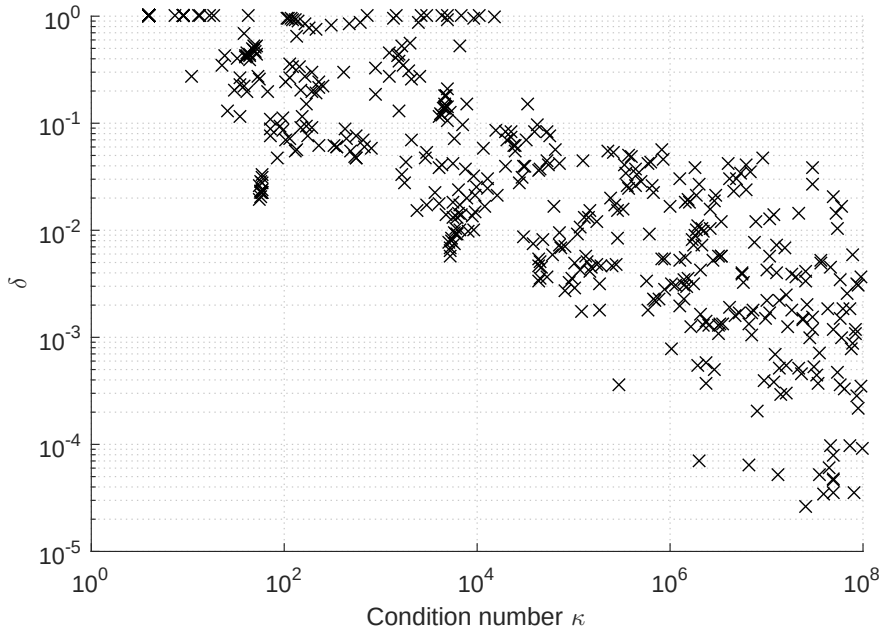

(B)

FIGURE 6. Iterations to $\epsilon = 10^{-6}$ residual convergence for the 1038 Newton direction problems described in-text: (a) ADMM vs GMRES-accelerated ADMM; (b) lower bounds on δ via Proposition 22.

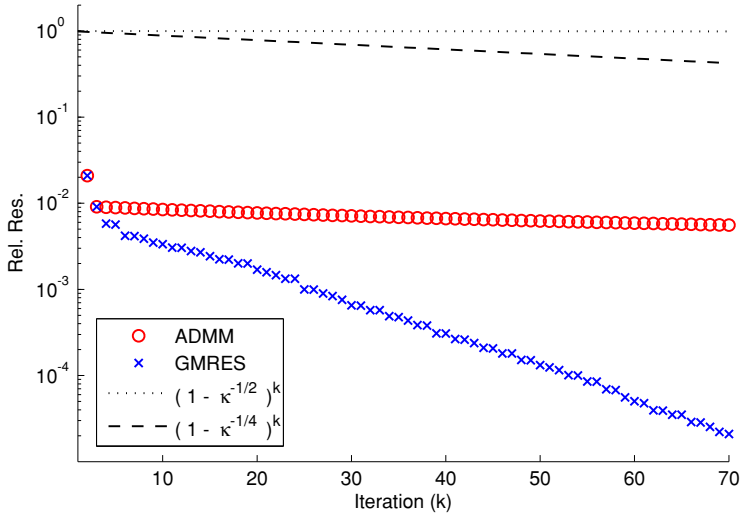

(A)

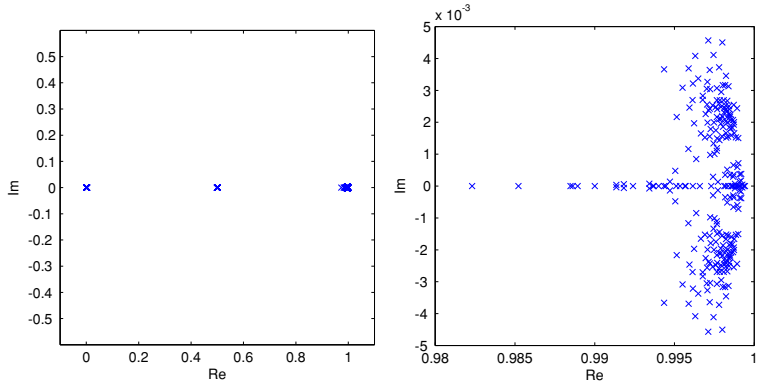

(B) (C)

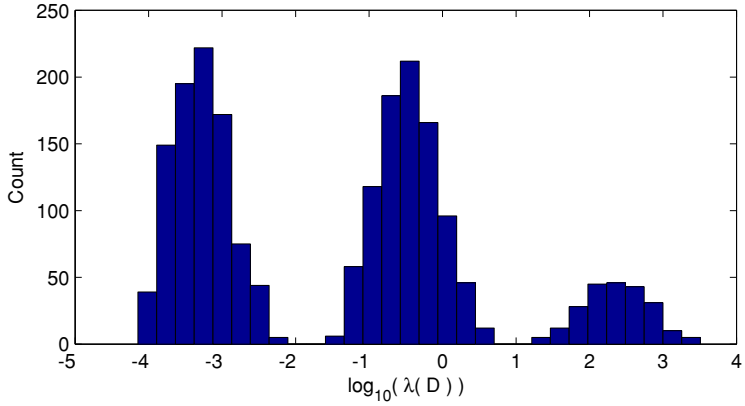

(D)

FIGURE 7. Plots for the 5th predictor step of "control3": (a) Convergence comparison; (b) Eigenvalues of G_{22}; (c) Eigenvalues of G_{22}, close-up around $+1$; (d) Histogram for eigenvalues of $\tilde{D}$.

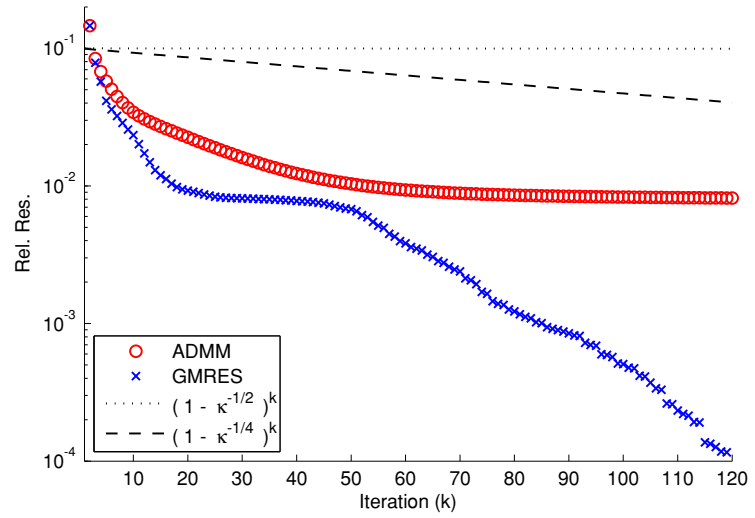

(A)

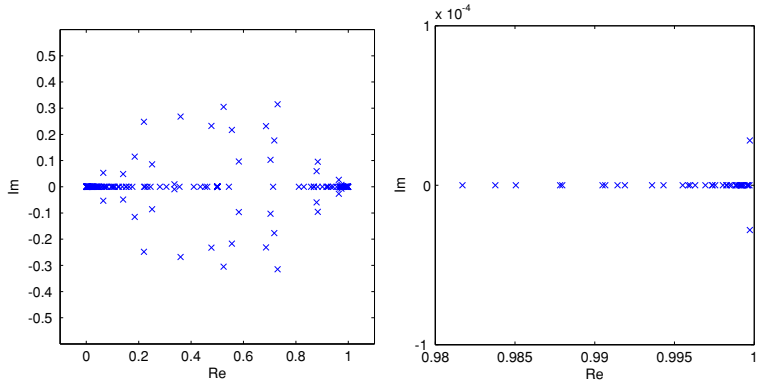

(B) (C)

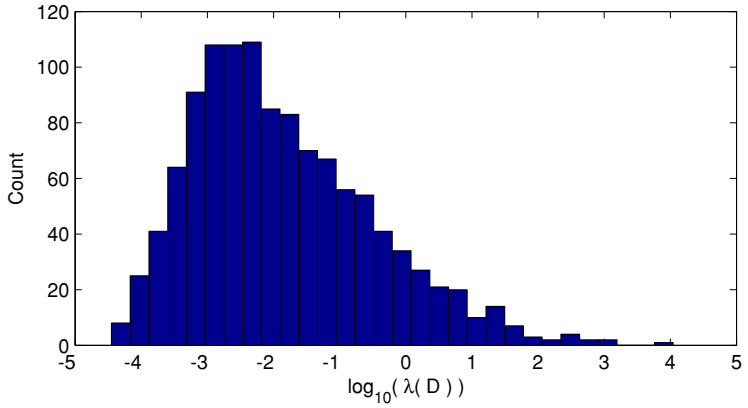

(D)

FIGURE 8. Plots for the 5th predictor step of "hinf14": (a) Convergence comparison; (b) Eigenvalues of G_{22}; (c) Eigenvalues of G_{22}, close-up around $+1$; (d) Histogram for eigenvalues of $\tilde{D}$.